

EXPRESS MAIL MAILING LABEL NUMBER EV318622542

PATENT

10830.0102.NPUS00

APPLICATION FOR UNITED STATES LETTERS PATENT

for

INTERNET PROTOCOL BASED DISASTER RECOVERY OF A SERVER

By

**John M. Hayden
Hongmei Wang
Frederic Corniquet
Philippe Armangau
Pascal Donette
Aju John**

BACKGROUND OF THE INVENTION

1. Field of the Invention

[0001] The present invention relates generally to data processing networks, and more particularly to disaster recovery of servers in a data processing network.

3. Description of Related Art

[0002] Large institutional users of data processing systems such as banks, insurance companies, and stock market traders must and do take tremendous steps to insure back up data availability in case of a major disaster. More recently these institutions have developed a heightened awareness of the importance of reducing business interruptions due to server failure in view of the many natural disasters and other world events including power grid failures and the destruction of the World Trade Center in New York City.

[0003] Business interruption due to loss of a server at a primary site can be minimized by maintaining a replica of the server at a secondary site. The state of the server at the secondary site should be synchronized to a prior consistent state of the server at the primary site in order to reduce disruption when service is switched over to the secondary site upon loss of service at the primary site. This has been done by remote data mirroring of storage volumes or files in the primary site to the secondary site.

[0004] It is possible to minimize disruption due to server failure by synchronous mirroring of server state and stored data between a primary server and a secondary server. This technique is practical when the primary server and the secondary server are close together so that there is no more than a tolerable loss of performance due to the delay in transmitting state and data changes from the primary server to the

1 secondary server and returning an acknowledgement from the secondary server to the
2 primary server.

3 [0005] Unfortunately, for disaster protection, the primary server and
4 secondary server must be geographically separated. For protection from large-scale
5 disasters such as hurricanes, power grid failure, and nuclear attack, the primary and
6 secondary servers should be separated by hundreds of kilometers. Once the primary and
7 secondary servers are separated by hundreds of kilometers, it is impractical to mirror
8 synchronously all of the state and data between the primary and secondary servers. This
9 is especially true when it is desired to use a shared data transmission link rather than a
10 dedicated link between the primary and secondary servers.

11 [0006] For remote mirroring, it is often impractical for the primary and
12 secondary servers to have identical software or hardware. In the usual case, the primary
13 server has been installed for quite some time in a local data network and the hardware
14 and software of the primary server is obsolete. Yet there is no desire to replace the
15 existing primary server because of the cost of labor and possible disruption of service
16 during replacement. In this situation, it is often possible to use a single present
17 generation hardware platform to support multiple virtual secondary servers.

18 [0007] A virtual secondary server provides replicas of the files stored on the
19 primary server. The files include user data files, server configuration files, and user
20 configuration information (such as user accounts, file ownership, and trustee rights) so
21 that users may login immediately and access data in the event of a failure of the primary
22 server.

1 **[0008]** For the typical case where the primary server and the virtual secondary
2 server are supported by different hardware and software, there have been problems
3 setting up the virtual secondary server in such a way that users can immediately access
4 data in the event of a failure on the primary server. The hardware supporting the virtual
5 secondary server must have sufficient storage and network resource for the virtual
6 secondary server. The network resources must be set up to avoid address conflicts.
7 Problems with storage resources, network resources, and address conflicts may not
8 become known until an attempted failover from the primary server to the virtual
9 secondary server.

11 **SUMMARY OF THE INVENTION**

12 **[0009]** In accordance with a first aspect, the invention provides a method in a
13 disaster recovery environment including a primary file server at an active site and a
14 secondary virtual file server at a disaster recovery site remote from the active site. The
15 secondary virtual file server includes a collection of files being replicated from the
16 primary server to the disaster recovery site. The secondary virtual file server needs
17 resources including network interfaces and file system mounts at the disaster recovery
18 site for providing user access at the disaster recovery site. The method includes
19 determining whether there are sufficient network interfaces and file system mounts at the
20 disaster recovery site for the virtual secondary file server for providing user access at the
21 disaster recovery site, and upon finding that there are sufficient network interfaces and
22 file system mounts at the disaster recovery site for the virtual secondary file server for
23 providing user access at the disaster recovery site, reserving the network interfaces and

1 file system mounts that are needed at the disaster recovery site for providing user access
2 at the disaster recovery site.

3 [00010] In accordance with another aspect, the invention provides a method in
4 a disaster recovery environment including a primary file server at an active site and a
5 secondary virtual file server at a disaster recovery site remote from the active site. The
6 secondary virtual file server includes a collection of files being replicated from the
7 primary server to the disaster recovery site. The secondary virtual file server needs
8 resources including network interfaces and file system mounts at the disaster recovery
9 site for providing user access at the disaster recovery site. The method includes
10 determining whether there are sufficient network interfaces and file system mounts at the
11 disaster recovery site for the virtual secondary file server for providing unrestricted user
12 access at the disaster recovery site once a configuration change would be made to the
13 primary file server, and upon finding that there are insufficient network interfaces and file
14 system mounts at the disaster recovery site for the virtual secondary file server for
15 providing unrestricted user access at the disaster recovery site once the configuration
16 change would be made to the primary file server, providing an operator with a list of
17 missing resources or discrepancies, and receiving from the operator a choice of
18 termination or configuration change, and upon receiving from the operator a choice of
19 configuration change, reserving network interfaces and file system mounts that are
20 available and needed at the disaster recovery site for providing user access at the disaster
21 recovery site, and then performing failover of user access from the active site to the
22 disaster recovery site.

1 **[00011]** In accordance with another aspect, the invention provides a method in
2 a disaster recovery environment including a primary file server at an active site and a
3 secondary virtual file server at a disaster recovery site remote from the active site. The
4 secondary virtual file server includes a collection of files being replicated from the
5 primary server to the disaster recovery site. The secondary virtual file server needs
6 resources including network interfaces and file system mounts at the disaster recovery
7 site for providing user access at the disaster recovery site. The method includes
8 determining whether there are sufficient network interfaces and file system mounts at the
9 disaster recovery site for the virtual secondary file server for providing unrestricted user
10 access at the disaster recovery site, and upon finding that there are insufficient network
11 interfaces and file system mounts at the disaster recovery site for the virtual secondary
12 file server for providing unrestricted user access at the disaster recovery site, providing an
13 operator with a list of missing resources or discrepancies, and receiving from the operator
14 a choice of termination or forced failover, and upon receiving from the operator a choice
15 of forced failover, reserving network interfaces and file system mounts that are available
16 and needed at the disaster recovery site for providing user access at the disaster recovery
17 site, and then performing failover of user access from the active site to the disaster
18 recovery site.

19 **[00012]** In accordance with another aspect, the invention provides a method in
20 a disaster recovery environment including a primary file server at an active site and a
21 secondary virtual file server at a disaster recovery site remote from the active site. The
22 secondary virtual file server includes a collection of files being replicated from the
23 primary server to the disaster recovery site. The method includes maintaining a primary

1 copy of user mappings at the disaster recovery site and a read-only cache of the user
2 mappings at the active site during user file access at the active site, and upon failover of
3 user access from the primary file server at the active site to the virtual secondary server at
4 the disaster recovery site, accessing the primary copy of user mappings at the disaster
5 recovery site in order to continue user file access at the disaster recovery site.

6 [00013] In accordance with another aspect, the invention provides a method in
7 a disaster recovery environment including a primary file server at an active site and a
8 secondary virtual file server at a disaster recovery site remote from the active site. The
9 secondary virtual file server includes a collection of files being replicated from the
10 primary server to the disaster recovery site. The method includes maintaining a copy of
11 user session information at the disaster recovery site during user file access at the active
12 site, and upon failover of user access from the primary file server at the active site to the
13 virtual secondary server at the disaster recovery site, accessing the copy of the user
14 session information at the disaster recovery site so that users accessing files of the
15 primary file server at the active site continue to access copies of the files at the disaster
16 recovery site without a need to log onto the disaster recovery site.

17 [00014] In accordance with yet another aspect, the invention provides a method
18 in a disaster recovery environment including a primary file server at an active site and a
19 secondary virtual file server at a disaster recovery site remote from the active site. The
20 secondary virtual file server includes a collection of files being replicated from the
21 primary server to the disaster recovery site. The method includes a network client
22 accessing the primary file server at the active site using a CIFS connection and detecting
23 a failure of the primary file server to respond to a file access request in a timely fashion,

1 and upon detecting the failure of the primary file server to respond to the file access
2 request in a timely fashion, the network client redirecting the file access request to the
3 disaster recovery site without terminating the CIFS connection.

4 **[00015]** In accordance with yet another aspect, the invention provides a method
5 in a disaster recovery environment including a primary file server at an active site and a
6 secondary virtual file server at a disaster recovery site remote from the active site. The
7 primary file server stores a collection of user files, and the secondary virtual file server
8 stores secondary copies of the user files. The method includes replicating changes to the
9 user files from the primary file server to the secondary copies of the user files in the
10 secondary virtual file server during user file access at the active site, and during the
11 replication of the changes to the user files from the primary file server to the secondary
12 virtual file server, creating at the disaster recovery site a series of snapshot copies of the
13 secondary virtual file server, each of the snapshot copies providing a group consistent
14 state of the user files in the secondary virtual file server.

15 **[00016]** In accordance with still another aspect, the invention provides a
16 disaster recovery system including a primary file server at an active site, and a secondary
17 virtual file server at a disaster recovery site remote from the active site. The secondary
18 virtual file server includes a collection of files being replicated from the primary server to
19 the disaster recovery site. The secondary virtual file server needs resources including
20 network interfaces and file system mounts at the disaster recovery site for providing user
21 access at the disaster recovery site. The disaster recovery system is programmed for
22 responding to a request from a system administrator by determining whether there are
23 sufficient network interfaces and file system mounts at the disaster recovery site for the

1 virtual secondary file server for providing user access at the disaster recovery site, and
2 upon finding that there are sufficient network interfaces and file system mounts at the
3 disaster recovery site for the virtual secondary file server for providing user access at the
4 disaster recovery site, reserving the network interfaces and file system mounts that are
5 needed at the disaster recovery site for providing user access at the disaster recovery site.

6 [00017] In accordance with yet another aspect, the invention provides a
7 disaster recovery system including a primary file server at an active site, and a secondary
8 virtual file server at a disaster recovery site remote from the active site. The secondary
9 virtual file server includes a collection of files being replicated from the primary server to
10 the disaster recovery site. The secondary virtual file server needs resources including
11 network interfaces and file system mounts at the disaster recovery site for providing user
12 access at the disaster recovery site. The disaster recovery system is programmed for
13 responding to a configuration change request from a system administrator by determining
14 whether there are sufficient network interfaces and file system mounts at the disaster
15 recovery site for the virtual secondary file server for providing unrestricted user access at
16 the disaster recovery site once a configuration change would be made to the primary file
17 server, and upon finding that there are insufficient network interfaces and file system
18 mounts at the disaster recovery site for the virtual secondary file server for providing
19 unrestricted user access at the disaster recovery site once a configuration change would
20 be made to the primary file server, providing the system administrator with a list of
21 missing resources or discrepancies, and receiving from the system administrator a choice
22 of termination or configuration change, and upon receiving from the system administrator
23 a choice of configuration change, reserving network interfaces and file system mounts

1 that are available and needed at the disaster recovery site for providing user access at the
2 disaster recovery site once the configuration change would be made to the primary file
3 server, and then performing the configuration change to the primary file server.

4 [00018] In accordance with yet another aspect, the invention provides a
5 disaster recovery system including a primary file server at an active site, and a secondary
6 virtual file server at a disaster recovery site remote from the active site. The secondary
7 virtual file server includes a collection of files being replicated from the primary server to
8 the disaster recovery site. The secondary virtual file server needs resources including
9 network interfaces and file system mounts at the disaster recovery site for providing user
10 access at the disaster recovery site. The disaster recovery system is programmed for
11 responding to a failover request from a system administrator by determining whether
12 there are sufficient network interfaces and file system mounts at the disaster recovery site
13 for the virtual secondary file server for providing unrestricted user access at the disaster
14 recovery site, and upon finding that there are insufficient network interfaces and file
15 system mounts at the disaster recovery site for the virtual secondary file server for
16 providing unrestricted user access at the disaster recovery site, providing the system
17 administrator with a list of missing resources or discrepancies, and receiving from the
18 system administrator a choice of termination or forced failover, and upon receiving from
19 the system administrator a choice of forced failover, reserving network interfaces and file
20 system mounts that are available and needed at the disaster recovery site for providing
21 user access at the disaster recovery site, and then performing failover of user access from
22 the active site to the disaster recovery site.

1 **[00019]** In accordance with yet another aspect, the invention provides a
2 disaster recovery system including a primary file server at an active site, and a secondary
3 virtual file server at a disaster recovery site remote from the active site. The secondary
4 virtual file server includes a collection of files being replicated from the primary server to
5 the disaster recovery site. The disaster recovery system is programmed for maintaining a
6 primary copy of user mappings at the disaster recovery site and a read-only cache of the
7 user mappings at the active site during user file access at the active site, and upon failover
8 of user access from the primary file server at the active site to the virtual secondary server
9 at the disaster recovery site, for accessing the primary copy of user mappings at the
10 disaster recovery site in order to continue user file access at the disaster recovery site.

11 **[00020]** In accordance with yet still another aspect, the invention provides a
12 disaster recovery system including a primary file server at an active site, and a secondary
13 virtual file server at a disaster recovery site remote from the active site. The secondary
14 virtual file server includes a collection of files being replicated from the primary server to
15 the disaster recovery site. The disaster recovery site is programmed for maintaining a
16 copy of user session information at the disaster recovery site during user file access at the
17 active site, and upon failover of user access from the primary file server at the active site
18 to the virtual secondary server at the disaster recovery site, accessing the copy of the user
19 session information at the disaster recovery site so that users accessing files of the
20 primary file server at the active site continue to access copies of the files at the disaster
21 recovery site without a need to log onto the disaster recovery site.

22 **[00021]** In accordance with a final aspect, the invention provides a disaster
23 recovery system including a primary file server at an active site, and a secondary virtual

1 file server at a disaster recovery site remote from the active site. The secondary virtual
2 file server includes a collection of files being replicated from the primary server to the
3 disaster recovery site. The disaster recovery system further includes at least one network
4 client programmed for accessing the primary file server at the active site using a CIFS
5 connection and detecting a failure of the primary file server to respond to a file access
6 request in a timely fashion, and upon detecting the failure of the primary file server to
7 respond to the file access request in a timely fashion, redirecting the file access request to
8 the disaster recovery site without terminating the CIFS connection.

9 [00022] In accordance with a final aspect, the invention provides a disaster
10 recovery system including a primary file server at an active site, and a secondary virtual
11 file server at a disaster recovery site remote from the active site. The primary file server
12 stores a collection of user files, and the secondary virtual file stores secondary copies of
13 the user files. The system is programmed for replicating changes to the user files from
14 the primary file server to the secondary copies of the user files in the secondary virtual
15 file server during user file access at the active site. The disaster recovery site is
16 programmed for creating at the disaster recovery site a series of snapshot copies of the
17 secondary virtual file server during the replication of the changes to the user files from
18 the primary file server to the secondary virtual file server. Each of the snapshot copies
19 provides a group consistent state of the user files in the secondary virtual file server.

20 21 BRIEF DESCRIPTION OF THE DRAWINGS

22 [00023] Other objects and advantages of the invention will become apparent
23 upon reading the detailed description with reference to the drawings, in which:

1 **[00024]** FIG. 1 is a block diagram of an IP network including multiple hosts
2 and multiple data mover computers;

3 **[00025]** FIG. 2 is a block diagram showing a primary data mover distributing
4 remote copy data to multiple secondary data movers in the IP network by establishing a
5 Transmission Control Protocol (TCP) connection with each of the secondary data
6 movers;

7 **[00026]** FIG. 3 is a block diagram showing a primary data mover distributing
8 remote copy data to multiple data movers through forwarder data movers;

9 **[00027]** FIG. 4 is a block diagram showing a shared save volume used to buffer
10 local copy data transmitted from a primary data mover to a secondary data mover;

11 **[00028]** FIG. 5 is a block diagram showing a primary save volume and a
12 secondary save volume;

13 **[00029]** FIG. 6 is a flowchart showing local replication in the system of FIG. 4;

14 **[00030]** FIG. 7 is a flowchart showing remote replication in the system of FIG.
15 5;

16 **[00031]** FIG. 8 is a block diagram of a primary site, including layered
17 programming in a primary data mover;

18 **[00032]** FIG. 9 is a block diagram of a secondary site, including layered
19 programming in a secondary data mover;

20 **[00033]** FIG. 10 is a flowchart of a process of replication at the primary site of
21 FIG. 8;

1 **[00034]** FIG. 11 is a flowchart of a procedure for producing a new remote copy
2 of a primary file system concurrent with ongoing replication and multicasting of
3 modifications to the primary file system;

4 **[00035]** FIG. 12 is a flowchart of an IP-replication send-thread introduced in
5 FIG. 8;

6 **[00036]** FIG. 13 is a block diagram of a volume multicast level in the data
7 mover programming of FIG. 8 and FIG. 9;

8 **[00037]** FIG. 14 is a block diagram of a replication control protocol (RCP)
9 level in the primary data mover programming of FIG. 8;

10 **[00038]** FIG. 15 is a block diagram of the RCP level in the secondary data
11 mover programming of FIG. 9;

12 **[00039]** FIG. 16 is a block diagram of a RCP forwarder at the RCP level in a
13 forwarder data mover;

14 **[00040]** FIG. 17 is a flowchart of an inbound RCP session in the secondary
15 data mover;

16 **[00041]** FIG. 18 is a block diagram showing a forwarder data mover
17 performing local replication;

18 **[00042]** FIG. 19 is a block diagram showing the sharing of a data mover's
19 single TCP port for RCP- connections with Hypertext Transfer Protocol (HTTP)
20 connections;

21 **[00043]** FIG. 20 is a block diagram showing the state of a remote replication
22 system before a disaster at the primary site;

1 **[00044]** FIG. 21 is a flowchart of a failover and recovery method performed
2 upon the remote replication system of FIG. 20 in order to recover from the disaster at the
3 primary site;

4 **[00045]** FIG. 22 is a block diagram showing the state of the remote replication
5 system of FIG. 20 after failover from the primary site to the secondary site;

6 **[00046]** FIG. 23 is a flowchart of a subroutine of failover with a checkpoint
7 and without sync, as used in the flowchart of FIG. 21;

8 **[00047]** FIG. 24 shows a block map that can be used to create a snapshot copy
9 of the secondary file system at a restart point during the failover of FIG. 23;

10 **[00048]** FIG. 25 shows a snapshot copy facility that can be used to create the
11 snapshot copy of the secondary file system at the restart point during the failover of FIG.
12 23;

13 **[00049]** FIG. 26 is a flowchart of a procedure for writing a block of data to a
14 production file system in the snapshot copy facility of FIG. 25;

15 **[00050]** FIG. 27 is a flowchart of a procedure for reading a block of data from
16 a snapshot file system in the snapshot copy facility of FIG. 25;

17 **[00051]** FIG. 28 is a flowchart of a subroutine to resync the primary file system
18 with the secondary file system, as used in the flowchart of FIG. 21;

19 **[00052]** FIG. 29 shows the state of the remote replication system of FIG. 20
20 during the resync procedure;

21 **[00053]** FIG. 30 is a flowchart of a subroutine to failback to the primary file
22 system, as used in the flowchart of FIG. 21;

1 **[00054]** FIG. 31 is a flowchart of execution of a failover command having a
2 sync option and a checkpoint option;

3 **[00055]** FIG. 32 is a flowchart of a subroutine for failover without sync, as
4 used in the flowchart of FIG. 31;

5 **[00056]** FIG. 33 is a subroutine for failover with sync, as used in the flowchart
6 of FIG. 31;

7 **[00057]** FIGS. 34 and 35 comprise a flowchart of a procedure for copying
8 snapshot delta for snapshots L to M;

9 **[00058]** FIG. 36 shows a block diagram of a data network in which snapshot
10 deltas are transmitted over a wide-area network from a network file server to a local file
11 server in order to update the local file system as needed;

12 **[00059]** FIGS. 37 and 38 comprise a flowchart of a procedure for replicating
13 snapshots in the data network of FIG. 36;

14 **[00060]** FIG. 39 is a modified version of the flowchart of FIG. 34, showing
15 how to copy a snapshot delta from a snapshot copy facility that keeps track of blocks of a
16 production file system that are not in use in the snapshot copies of the production file
17 system;

18 **[00061]** FIG. 40 shows a specific instance of the flowchart of FIG. 39 for a
19 snapshot copy facility that uses respective meta bit maps for indicating the blocks in a
20 production file system that are not used in the snapshot copies of the production file
21 system;

1 **[00062]** FIG. 41 is a flowchart for writing a specified data block to the
2 production file system for a snapshot copy facility that uses a meta bit map for indicating
3 the blocks in a production file system that are not presently in use;

4 **[00063]** FIG. 42 is a diagram of a preferred organization of multiple snapshots
5 in the snapshot copy facility;

6 **[00064]** FIG. 43 shows a bit map including a page table and a set of pages, for
7 use in the snapshot organization of FIG. 42;

8 **[00065]** FIG. 44 shows a block map including a hash table and hash lists, for
9 use in the snapshot organization of FIG. 42;

10 **[00066]** FIG. 45 shows a specific construction for and interpretation of a meta
11 bit map for the production volume; and

12 **[00067]** FIG. 46 shows an alternative interpretation of a meta bit map for the
13 production volume.

14 **[00068]** FIG. 47 shows virtual data movers and virtual servers in the network
15 file server introduced in FIG. 1;

16 **[00069]** FIG. 48 shows a conventional configuration of a data mover and a file
17 system for the data mover;

18 **[00070]** FIG. 49 shows the installation of two virtual data movers in the data
19 mover of FIG. 48;

20 **[00071]** FIG. 50 shows various components of a virtual data mover;

21 **[00072]** FIG. 51 is a flowchart of a procedure for configuring a virtual data
22 mover;

1 **[00073]** FIG. 52 shows a menu of operations selectable by a system
2 administrator for management of a virtual data mover;

3 **[00074]** FIG. 53 is a flowchart of a procedure for prevalidation and resource
4 reservation for a configuration change of a virtual data mover or for failover of a virtual
5 data mover from an active site to a disaster recovery site ;

6 **[00075]** FIG. 54 shows user mapping facilities in a primary data mover at an
7 active site and a secondary data mover at a disaster recovery site;

8 **[00076]** FIG. 55 is a flowchart of a procedure for using the user mapping
9 facilities of FIG. 54;

10 **[00077]** FIG. 56 is a block diagram showing a number of agents in a disaster
11 recovery environment for maintaining a client-server connection during a disaster;

12 **[00078]** FIG. 57 is a flowchart of a procedure used in the disaster recovery
13 environment of FIG. 56 for maintaining the client-server connection;

14 **[00079]** FIG. 58 is a block diagram showing a series of snapshot copies of a
15 virtual data mover at the disaster recovery site; and

16 **[00080]** FIG. 59 is a flowchart of a procedure used at the disaster site for
17 creating the series of snapshot copies of the virtual data mover at the disaster recovery
18 site.

19 **[00081]** While the invention is susceptible to various modifications and
20 alternative forms, specific embodiments thereof have been shown by way of example in
21 the drawings and will be described in detail. It should be understood, however, that it is
22 not intended to limit the form of the invention to the particular forms shown, but on the

contrary, the intention is to cover all modifications, equivalents, and alternatives falling within the scope of the invention as defined by the appended claims.

DESCRIPTION OF ILLUSTRATIVE EMBODIMENTS

Remote Data Replication Facility

[00082] FIG. 1 shows an IP network 220 including multiple network file servers 221, 222, and multiple hosts 223, 224, 225. The hosts and network file servers, for example, can be distributed world wide and linked via the Internet. Each of the network file servers 221, 222, for example, has multiple data movers 226, 227, 228, 232, 233, 234, for moving data between the IP network 220 and the cached disk array 229, 235, and a control station 230, 236 connected via a dedicated dual-redundant data link 231, 237 among the data movers for configuring the data movers and the cached disk array 229, 235. Further details regarding the network file servers 221, 222 are found in Vahalia et al., U.S. Patent 5,893,140, incorporated herein by reference.

[00083] In operation, it is desired for each of the network file servers 221, 222 to provide read-only access to a copy of the same file system. For example, each of the network file servers could be programmed to respond to user requests to access the same Internet site. The IP network 220 routes user requests to the network file servers 221, 222 in the same continent or geographic region as the user. In this fashion, the user load is shared among the network file servers.

[00084] In the wide-area network of FIG. 1, it is desired to perform read-write updating of the respective file system copies in the network file servers 221, 222 while

1 permitting concurrent read-only access by the hosts. It is also desired to distribute the
2 updates over the IP network.

3 [00085] There are a number of ways that updates could be distributed over the
4 IP network from a primary data mover to multiple secondary data movers. As shown in
5 FIG. 2, for example, a primary data mover establishes a connection 242, 243, 244 in
6 accordance with the industry-standard Transmission Control Protocol (TCP) over the IP
7 network 220 to each secondary data mover 245, 246, 247, and then concurrently sends
8 the updates to each secondary data mover over the TCP connection. When the updates
9 need to be distributed to a large number of secondary data movers, however, the amount
10 of time for distributing the updates may become excessive due to limited resources (CPU
11 execution cycles, connection state, or bandwidth) of the primary data mover 241. One
12 way of extending these limited resources would be to use existing IP routers and switches
13 to implement “fan out” from the primary data mover 241 to the secondary data movers
14 245, 246, 247. Still, a mechanism for reliability should be layered over the Internet
15 Protocol.

16 [00086] FIG. 3 shows that the time for distributing updates from a primary data
17 mover 251 to a large number of secondary data movers 254, 255, 256, 257 can be
18 reduced by using intermediate data movers 252, 253 as forwarders. The primary data
19 mover 251 sends the updates to the forwarder data movers 252, 253, and each of the
20 forwarder data movers sends the updates to a respective number of secondary data
21 movers. The forwarder data movers 252, 253 may themselves be secondary data movers;
22 in other words, each may apply the updates to its own copy of the replicated read-only
23 file system. The distribution from the primary data mover 251 to the forwarder data

1 movers 252, 253 can be done in a fashion suitable for wide-area distribution (such as over
2 TCP connections). The forwarding method of replication of FIG. 3 also has the
3 advantage that the distribution from each forwarder data mover to its respective data
4 movers can be done in a different way most suitable for a local area or region of the
5 network. For example, some of the forwarder data movers could use TCP connections,
6 and others could use a combination of TCP connections for control and UDP for data
7 transmission, and still other forwarders could be connected to their secondary data
8 movers by a dedicated local area network.

9 [00087] For implementing the replication method of FIG. 3 over the Internet
10 Protocol, there are a number of desired attributes. It is desired to maintain independence
11 between the primary data mover and each of the secondary data movers. For example, a
12 new secondary data mover can be added at any time to replicate an additional remote
13 copy. The primary data mover should continue to function even if a secondary data
14 mover becomes inoperative. It is also desired to maintain independence between the
15 replication method and the IP transport method. Replication should continue to run even
16 if the IP transport is temporarily inactive. It is desired to recover in a consistent fashion
17 from a panic or shutdown and reboot. A record or log of the progress of the replication
18 can be stored for recovery after an interruption. It is desired to build re-usable program
19 blocks for the replication function, so that the program blocks for the replication function
20 can be used independent of the location of the primary file system or its replicas.

21 [00088] In a preferred implementation, independence between the replication
22 process, the IP transport method, and the primary file system being replicated, is ensured
23 by use of a save volume. The save volume is a buffer between the data producer (i.e., the

1 host or application updating the primary file system), the replication process, and the data
2 consumer (the secondary data movers). The save volume stores the progress of the
3 replication over the Internet Protocol so as to maintain the consistency of the replication
4 process upon panic, reboot, and recovery. The transport process need not depend on any
5 “in memory” replication information other than the information in the save volume, so as
6 to permit the replication process to be started or terminated easily on any data mover for
7 load shifting or load balancing.

8 [00089] When a save volume is used, it can be shared between a primary data
9 mover and a secondary data mover in the case of local file system replication, or a
10 primary copy of the shared volume can be kept at the primary site, and a secondary copy
11 of the shared volume can be kept at the secondary site, in the case of remote file system
12 replication.

13 [00090] For the case of local file system replication, FIG. 4 shows a primary
14 site including a primary data mover 260 managing access to a primary file system 261,
15 and a secondary data mover 262 managing access to a secondary file system 263
16 maintained as a read-only copy of the primary file system 261. A save volume 264 is
17 shared between the primary data mover 260 and the secondary data mover 262. This
18 sharing is practical when the secondary site is relatively close to the primary site. A redo
19 log 265 records a log of modifications to the primary file system 261 during the
20 replication process for additional protection from an interruption that would require a
21 reboot and recovery.

22 [00091] Local replication can be used to replicate files within the same
23 network file server. For example, in the network file server 221 in FIG. 1, the primary

1 data mover could be the data mover 226, the secondary data mover could be the data
2 mover 227, the save volume could be stored in the cached disk array 229, and replication
3 control messages could be transmitted between the data movers over the data link 231.

4 **[00092]** For the case of remote file system replication, FIG. 5 shows a primary
5 site including a primary data mover 270 managing access to a primary file system 271,
6 and a secondary data mover 272 managing access to a secondary file system 273
7 maintained as a read-only copy of the primary file system 271. The primary site includes
8 a primary save volume 274, and the remote site includes a secondary save volume 275.
9 A redo log 276 records a log of modifications to the primary file system 271 during the
10 replication process for additional protection from an interruption that would require a
11 reboot and recovery.

12 **[00093]** FIG. 6 shows a method of operating the system of FIG. 4 for local
13 replication. In a first step 281, the primary data mover migrates a copy of the primary
14 file system to create a secondary file system at the secondary site in such a way to permit
15 concurrent write access to the primary file system. The migration, for example, may use
16 the method shown in FIG. 17 of Ofek U.S. Patent 5,901,327, in which a bit map indicates
17 remote write pending blocks. Alternatively, the migration may use a snapshot copy
18 mechanism, for example, as described in Kedem, U.S. Patent 6,076,148, in which a bit
19 map indicates the blocks that have changed since the time of snap-shotting of the primary
20 file system. The snapshot method is preferred, because it is most compatible with the
21 delta set technique for remote copy of subsequent modifications. For example, a
22 snapshot manager creates a snapshot copy of the primary file system, as will be further
23 described below with reference to FIGS. 25 to 27. The migration may involve taking a

1 first snapshot copy of the primary file system, copying the first snapshot copy to the
2 secondary file system, starting replication and creating a second snapshot copy of the
3 primary file system, and then copying the differential between the first snapshot copy and
4 the second snapshot copy from the primary file system to the secondary file system, in a
5 fashion similar to that described further below with reference to steps 536 and 537 in
6 FIG. 28. In any event, it is desired for the secondary file system to become a copy of the
7 state of the primary file system existing at some point of time, with any subsequent
8 modifications of the primary file system being transferred through the shared save
9 volume.

10 [00094] In step 282, the primary data mover writes subsequent modifications
11 of the primary file system to the shared save volume. In step 283, the secondary data
12 mover reads the subsequent modifications from the shared save volume and writes them
13 to the secondary file system. In step 284, the secondary data mover provides user read-
14 only access to consistent views of the secondary file system, by integrating the
15 subsequent revisions into the secondary file system while providing concurrent read-only
16 access to the secondary file system. This can be done by using a remote data facility to
17 maintain a remote copy of a pair of delta volumes and to switch between the delta
18 volumes as described in Suchitra Raman, et al., U.S. Patent Application Ser. No.
19 10/147,751 filed May 16, 2002, entitled "Replication of Remote Copy Data for Internet
20 Protocol (IP) transmission, Pub. No. US 2003/0217119 A1 published Nov. 20, 2003,
21 incorporated herein by reference. This can also be done by successively accumulating
22 delta sets at the secondary, and accessing a block index or map of updated data blocks in
23 the delta set.

1 **[00095]** Each delta set, for example, is a set of consistent updates to blocks of
2 the secondary file system. The updated blocks are included in “delta chunks” transmitted
3 from the primary file server to the secondary file server. Each delta set includes a
4 number of delta chunks, and each delta chunk belongs to a single delta set. Each delta
5 chunk is transmitted with a header including a generation count, a delta set number, and a
6 chunk number within the delta set. When a delta set is accumulated, a block index or
7 map is created to identify the updated blocks in the delta set. Once the delta set has been
8 accumulated, a background integration process is started that writes the updates into the
9 secondary file system, and the progress of the background process is recorded by
10 updating the block index or map for the delta set when each updated block is written to
11 the secondary file system. During the background process, read access to the secondary
12 file system is permitted on a priority basis. For this read access to a specified block of
13 the file system, the secondary file server first accesses the block index or map of the delta
14 set, and if the specified block is in the delta set, the secondary file server returns the data
15 of the block from the delta set. If the specified block is not in the delta set, then the
16 secondary file server returns the data of the block from the secondary file system.

17 **[00096]** In FIG. 6, execution loops from step 284 back to step 282. In this
18 fashion, the secondary file system is updated from the primary site concurrently with
19 read-only access at the secondary site.

20 **[00097]** FIG. 7 shows a method of operating the system of FIG. 5 for remote
21 replication. In a first step 291, the primary data mover migrates a copy of the primary
22 file system to create a secondary file system at the secondary site, in a fashion similar to
23 step 281 in FIG. 6. In step 292, the primary data mover writes subsequent modifications

1 of the primary file system to the primary save volume, in a fashion similar to step 282 in
2 FIG. 6. In step 293, the modifications are copied from the primary save volume to the
3 secondary save volume, by transmitting delta chunks. In step 294, the secondary data
4 mover reads the modifications from the secondary save volume and writes them to the
5 secondary file system. In step 295, the secondary data mover provides user read-only
6 access to consistent views of the secondary file system, in a fashion similar to step 284 of
7 FIG. 6. Execution loops from step 295 back to step 292. In this fashion, the secondary
8 file system is remotely updated from the primary site concurrently with read-only access
9 at the secondary site.

10 [00098] FIG. 8 shows layered programming 300 for a primary data mover. It
11 is desired to use layered programming in accordance with the International Standard
12 Organization's Open Systems Interconnection (ISO/OSI) model for networking protocols
13 and distributed applications. As is well known in the art, this OSI model defines seven
14 network layers, namely, the physical layer, the data link layer, the network layer, the
15 transport layer, the session layer, the presentation layer, and the application layer.

16 [00099] As shown in FIG. 8, the layered programming 300 includes a
17 conventional TCP/IP transport layer 301. The layers above the TCP/IP transport layer
18 301 include a replication control protocol (RCP) session layer 302, a volume multicast
19 presentation layer 303, and an IP-FS (file system) copy send-thread 304 and an IP-
20 replication send-thread 305 at the program layer level. Over these program layers is a
21 management and configuration command interpreter (MAC_CMD) 306 for system
22 operator set-up, initiation, and supervisory control of the replication process.

1 **[000100]** In operation, the RCP layer 302 provides an application program
2 interface (API) for multicasting data over TCP/IP. RCP provides callback,
3 acknowledgement (ACK), and resumption of aborted transfers.

4 **[000101]** RCP provides the capability for a remote site to replicate and
5 rebroadcast remote copy data. The remote site functions as a router when it rebroadcasts
6 the remote copy data. RCP can also be used to replicate data locally within a group of
7 data movers that share a data storage system.

8 **[000102]** To create a new remote copy in response to a supervisory command,
9 the command interpreter 306 initiates execution of a replication module 310 if the
10 replication module is not presently in an active mode. Then, the command interpreter
11 306 invokes a snapshot manager 308 to create a snapshot copy 309 of a primary file
12 system volume 307. When the snapshot copy is created, the snapshot manager 308
13 obtains a current delta set number from the replication module 310 and inserts the current
14 delta set number into the metadata of the snapshot. The current delta set number for the
15 snapshot is all that the secondary needs to identify modifications that are made
16 subsequent to the creation of the snapshot. In this fashion, any number of new remote
17 copies can be created at various times during operation of the replication module, with
18 the snapshot process operating concurrently and virtually independent of the replication
19 module. For example, whenever synchronization of a remote copy is lost, for example
20 due to a prolonged disruption of network traffic from the primary site to the remote site, a
21 new remote copy can be created to replace the unsynchronized remote copy.

22 **[000103]** Once the snapshot copy 309 is accessible, the command interpreter
23 306 initiates execution of an instance of the IP-FS copy send-thread 304. The instance of

1 the IP-FS copy send-thread 304 reads data from the snapshot copy 309 and calls upon the
2 volume multicast layer 303 to multicast the remote copy data to all of the secondary data
3 movers where the remote copies are to be created. This can be a copy by extent, so there
4 is no copying of invalid or unused data blocks. For example, the volume multicast layer
5 303 is given a copy command (@vol., length) specifying a volume and an extent to be
6 copied, and may also specify a group of destinations (an RCP group). The snapshot copy
7 309 of the primary file system identifies the next valid block to be copied, and the
8 number of valid contiguous blocks following the next block. These blocks are copied at
9 the logical level, so it does not matter what physical structure is used for storing the
10 secondary file system at the secondary site. The copying is done locally, or by remote
11 copy, for example by transporting the data block over IP. The volume multicast layer
12 303 invokes the RCP layer 302 to transport each data block.

13 [000104] During the remote copy process, whenever a modification is made to a
14 block of the primary file system volume 307, the replication module 310 logs an
15 indication of the modified block in a log 314 and later assembles the modification into a
16 delta set chunk written to a primary save volume 311. The replication module 310 logs
17 the indications in the log 314 on a priority or foreground basis as data is written to the
18 primary file system volume 307, and also logs boundaries between delta sets. The
19 replication module 310 later reads the log 314 to read the indicated modifications from
20 the primary file system volume 307, assemble the indicated modifications into delta set
21 chunks on a background basis, and store the delta set chunks in a save volume chunk area
22 of the save volume 311. For example, the log is in the form of a queue of two bit-map
23 tables, a new one of the tables being written to coincident with write operations upon the

1 primary file system volume 307, and an old one of the tables being read to determine
2 blocks to copy from the primary file system to create a new delta set in the save volume
3 311. When the delta set chunks become available for distribution from the save volume
4 311, the replication module 310 updates the save volume mailbox area 312 by storing
5 each delta set chunk definition (@vol., length).

6 [000105] The IP-replication send-thread instance 305 polls the save volume
7 mailbox area 312 to see if any delta set chunks have been stored in the save volume
8 chunk area 313. If so, then the thread instance calls upon the volume multicast layer 303
9 to multicast the delta set chunks to the data movers that manage the storage of the
10 respective remote file system copies. For example, for each delta set chunk, the IP-
11 replication send-thread instance 305 issues a volume multicast command to the volume
12 multicast layer 303. When the chunk multicast is completed, the IP-replication send-
13 thread instance 305 updates its context on the save volume 311 in the mailbox area 312.
14 At reboot after an interruption of multicast of a chunk, the IP-replication send-thread
15 instance is able to restart the multicast of the chunk. The IP-replication send-thread
16 instance also is responsible for retrying transmission of the chunk whenever the
17 connection with the secondary is interrupted.

18 [000106] FIG. 9 shows the layered programming 320 for a secondary data
19 mover. The programming includes a TCP/IP layer 321, an RCP layer 322, a volume
20 multicast layer 323, and a management and configuration command interpreter
21 (MAC_CMD) 324. During creation of a new remote copy in a secondary file system
22 volume 325, the volume multicast layer 323 writes remote copy data from the primary
23 data mover to the secondary file system volume 325, and concurrently writes

1 modifications (delta set chunks) from the primary data mover to a save volume chunk
2 area 326 of a secondary save volume 327.

3 [000107] A header for the changes in a next version of the delta set is sent last,
4 because there is no guarantee of the order of receipt of the IP packets. The header of the
5 delta set includes a generation count, the number of delta blocks for the next version of
6 the delta set, a checksum for the header, and a checksum for the data of all the delta
7 blocks. The receiver checks whether all of the changes indicated in the header have been
8 received.

9 [000108] Once a complete remote snapshot copy has been reconstructed in the
10 secondary file system volume 325, a playback module 328 is activated to read the
11 modifications from the save volume chunk area 326 and integrates them into the
12 secondary file system volume 325. From each delta-set chunk in the save volume area
13 326, the playback module 328 gets the block address and number of contiguous blocks to
14 be written to the secondary file system volume. An access module 329 provides read-
15 only access to a consistent view of the secondary file system in the secondary file system
16 volume 325.

17 [000109] FIG. 10 shows a procedure executed by the primary site of FIG. 8 to
18 perform replication of the primary file system. When replication is started in a first step
19 341, the primary file system is paused to make it consistent. Migration of the primary file
20 system to the secondaries can then be started using a remote copy facility or snapshot
21 manager. Then, in step 342, concurrent write access to the primary file system is
22 resumed, and all modifications made on the primary file system are logged at the volume

1 level on a priority or foreground basis when each modification is made. In addition, a
2 background process of delta-set creation is initiated.

3 **[000110]** Two configurable triggers specify the rate of delta set creation: a
4 timeout parameter and a high water mark parameter. Whenever delta set creation is
5 initiated, the current time, as indicated by a real-time clock, is added to a configurable
6 timeout interval to produce the timeout parameter. The high water mark specifies an
7 amount of modified data, in megabytes. The first trigger that occurs will trigger the
8 creation of a delta set. The replication module creates the delta set by pausing the
9 primary file system, copying the modified blocks from the primary file system to the
10 delta set volume, and then resuming the primary file system. By logging indications of
11 the modified blocks and later copying the modified blocks, multiple modifications to the
12 same block are represented and transported once during a single delta set.

13 **[000111]** In step 343, the background process of delta set creation is temporarily
14 suspended, for example, by placing the process on a task queue that is periodically
15 serviced. In step 344, execution of the delta set creation process is resumed. In step 345,
16 the modification size is compared to the high water mark. If the high water mark is not
17 exceeded, then execution continues to step 346. In step 346, the present value of the real-
18 time clock is compared to the timeout parameter. If the timeout parameter has not been
19 exceeded, then execution loops back to step 343. Otherwise, execution continues to step
20 347. Execution also branches to step 347 from step 345 if the modification size is greater
21 than the high water mark.

22 **[000112]** In step 347, the primary file system is paused. In step 348, a new delta
23 set is created by starting the copying of modified blocks from the primary file system

1 volume to the new delta set. In step 349, the logging of new modifications into a new
2 table is started. In step 350, the time-out and high water mark is re-armed. In other
3 words, a new value for the timeout parameter is computed as the current real time plus
4 the configurable timeout interval, and the modification size is reset to indicate the size of
5 the new modifications. In step 351, the primary file system is resumed. Execution loops
6 from step 351 back to step 343 to suspend the background process of delta set creation.

7 **[000113]** To maintain the consistency of the delta set created in the primary save
8 volume, the primary file system could remain paused and not resumed in step 351 until
9 the copy process begun in step 348 is completed. Preferably, however, the copy process
10 begun in step 348 is a snapshot copy process, so that write access to the primary file
11 system may resume in step 351 before the copy process has been completed. For the
12 example of the modification log being a queue of two bit-map tables, when a write access
13 to a block in the primary file system is requested, the old bit map is accessed on a priority
14 basis. If the corresponding bit in the old bit map indicates a modified block in the
15 primary file system volume not yet copied to the save volume, then it is copied on a
16 priority basis to the save volume before the new write data is written to the primary file
17 system volume. As soon as a modified block has been copied from the primary file
18 system volume to the save volume, the corresponding bit in the old bit map is cleared. In
19 this fashion, at the completion of the copy process, the entire old table will be in a reset
20 state, ready to be used as the next new table.

21 **[000114]** When the copy process started in step 348 is completed, the replication
22 module sets the save volume mailbox area to show that a new delta set is ready for
23 transmission. Upon polling the mailbox area, the IP-replication send-thread finds that the

1 new delta set is ready for transmission, and invokes the volume multicast layer to
2 transmit the delta set to the secondary sites. After step 351, execution loops back to step
3 343.

4 **[000115]** FIG. 11 shows a flow chart of the overall procedure of creating a new
5 remote copy, either for the first time at a secondary site or as a replacement for a remote
6 copy that needs to be resynchronized with the primary file system. In a first step 352, the
7 snapshot manager creates a snapshot copy of the primary file system at the end of any
8 pending transaction upon the primary file system (e.g., when the primary file system
9 becomes consistent after it is paused in step 341 of FIG. 10 or in step 347 of FIG. 10.)
10 The replication module independently writes any subsequent modifications into a current
11 delta set for the next transaction.

12 **[000116]** In step 353, the snapshot manager obtains the current delta set number
13 from the replication module and inserts it into metadata of the snapshot copy. In step
14 354, the IP-FS copy send-thread is started in order to send volume extents of the snapshot
15 copy to the secondary data mover, by invoking the volume multicast layer for each
16 extent.

17 **[000117]** In step 355, when the IP-FS copy send-thread is finished, the primary
18 data mover sends a “start playback” signal to the secondary data mover. In step 356, the
19 secondary data mover receives the “start playback” signal from the primary data mover,
20 and starts the playback module. In step 357, playback module begins playback from the
21 delta set indicated by the delta set number in the snapshot metadata.

22 **[000118]** The playback module (328 in FIG. 23) at the secondary site integrates
23 the delta set modifications into secondary file system. Each time that a new delta set

1 appears in the secondary save volume, the modifications can be integrated into the
2 secondary file system, for example, by pausing the secondary file system, copying the
3 modifications from the secondary save volume into the secondary file system, and
4 resuming the secondary file system. Alternatively, a timeout interval and a high water
5 mark value can be configured for the secondary site, so that the modifications may be
6 integrated into the secondary file system at a rate less frequent than the rate at which the
7 new delta sets appear in the secondary save volume. In this case, the modifications from
8 the secondary save volume would not be integrated into the secondary file system until
9 the timeout time is reached unless the amount of modifications in the save volume
10 reaches the high water mark. As described above, the integration of the modifications
11 can be performed concurrently with read-only access to a consistent view of the
12 secondary file system.

13 [000119] FIG. 12 shows a flowchart of the IP-replication send-thread (305 in
14 FIG. 8). In a first step 361, the thread polls the primary save volume mailbox area. If the
15 mailbox area indicates that there is not a new delta set chunk in the primary save volume
16 area, then the thread is finished for the present task invocation interval. Execution of the
17 thread is suspended in step 363, and resumed in step 364 at the next task invocation
18 interval.

19 [000120] If the mailbox area indicates that there is a new delta set chunk in the
20 primary save volume, then execution continues from step 362 to step 365. In step 365,
21 the IP-replication send-thread issues a volume multicast command to broadcast or
22 forward the delta set chunk to specified destination data movers. In step 366, if the
23 multicast has been successful, then execution branches to step 367. In step 367, the IP-

1 replication send-thread updates the primary save volume mailbox to indicate completion
2 of the multicast, and execution continues to step 363 to suspend execution of the thread
3 until the next task invocation interval.

4 **[000121]** In step 366, if the multicast is not successful, then execution continues
5 to step 368 to test whether more than a certain number (N) of retries have been attempted.
6 If not, then execution loops back to step 365 to retry the multicast of step 365. If more
7 than N retries have been attempted, then execution continues from step 368 to step 369.
8 In step 369, the IP-replication send-thread logs the error, and then in step 370, passes
9 execution to an error handler.

10 **[000122]** FIG. 13 shows various objects defined by the volume multicast layer.
11 The volume multicast layer provides multicast service to instances of a VolMCast object
12 370 representing a volume multicast relationship between a respective primary file
13 system volume specified by a volume name (volumeName) and a respective group of
14 secondary data movers specified by an RCP group name (rcpgpeName). For example, at
15 configuration time, one or more RCP groups are defined in response to configuration
16 commands such as:

17
18 .RCP_config <server_name> add <IP>
19

20 This configuration command adds the IP address (IP) of a specified destination data
21 mover (server_name) to an RCP group.

1 **[000123]** Also at configuration time, a specified data mover can be defined to be
2 a primary data mover with respect to the RCP group (a relationship called a
3 MultiCastNode) in response to a configuration command such as:

4
5 .server_config <server_name> rep <groupname> add <IP>

6
7 where “server_name” is the name for the primary data mover, “groupname” is the name
8 of a configured RCP group, and “IP” is the IP address of the primary data mover. When
9 configuration of the MultiCastNode object is finished, the MultiCastNode object will
10 have its own name, a name for the primary data mover, an RCP group name, and a list of
11 IP addresses to which the primary server should broadcast in order to transmit IP packets
12 to all the secondary data movers in the RCP group.

13 **[000124]** The VolMCast object can then be built on top of a MultiCastNode
14 object. The additional information required for the VolMCast object is, on the sender
15 side, the primary or source file system volume and on each receiver side, the secondary
16 or destination file system volume. For flexibility, it is permitted to specify a different
17 volume name on each secondary data mover. By specifying the destination volume
18 names during creation of the VolMCast object, it is not necessary to specify the
19 destination volume names at each copy time. For example, the VolMCast object is
20 defined by configuration commands to the primary data mover such as:

21
22 .server_config <server_name> "volmcast MultiCastNodeName>

23 [-src | -dest] volume"

1
2 where <server_name> is the name of the MultiCast Node.

3 **[000125]** Once the VolMCast object has been defined, an IP-replication service
4 can be configured for the object upon the primary data mover. Then the primary data
5 mover will respond to commands for starting the replication service and stopping the
6 replication service upon the VolMCast object. When replication is stopped on a
7 secondary, the secondary file system is left in a consistent state. In other words, if a
8 replay was in progress, the stop will complete when the replay is finished.

9 **[000126]** The primary data mover may respond to additional commands for
10 create a new delta set on demand, updating the replication policy (high water mark and
11 timeout interval parameters) on the primary file system or secondary file systems, and
12 defining persistency of the replication process upon remount or reboot of the primary file
13 system or any one of the secondary file systems. For example, at reboot the replication
14 service is re-started on the primary file system and the secondary file system in the state it
15 was at unmount or shutdown. A recovery of the replication context happens at reboot or
16 on remount. The replica recovery is executed before the primary and secondary file
17 systems are made available for user access. This allows all modifications during the
18 recovery of the primary file system to be logged by the replication service.

19 **[000127]** As shown in FIG. 13, the volume multicast layer is responsive to a
20 number of commands 371 from higher layers in the protocol stack. In addition to the
21 configuration commands for defining a new VolMCast object relating a specified primary
22 file system volume to a specified RCP group, an existing VolMCast object can be opened
23 for either a sender mode or a receiver mode. An opened VolMCast object can be closed.

1 Once a VolMCast object has been opened in a sender mode, it can be called upon to
2 broadcast a control block (CB) to the secondary volumes of the VolMCast object, such as
3 a control block specifying a remote copy of a specified extent of the primary volume.

4 **[000128]** Control blocks may specify various operations upon the secondary
5 volumes of the VolMCast object, such as cluster file system commands for performing
6 operations such as invalidations, deletions, renaming, or other changes in the
7 configuration of the objects of the file system upon all copies (local or remote) of the file
8 system. In this case, RCP is used for the broadcast or forwarding of the cluster file
9 system commands to all the data movers that are to operate upon the local or remote
10 copies of the file system, and for returning acknowledgement of completion of the
11 operations upon all of the copies of the file system.

12 **[000129]** With reference to FIG. 13, the volume multicast layer defines a
13 VolMCastSender object 372 instantiated when a VolMCast instance is opened in the
14 sending mode, and a VolMCastReceiver object 373 instantiated when a VolMCast
15 instance is opened in a receiving mode. The VolMCastSender object class and the
16 VolMCastReceiver object class inherit properties of the VolMCast object class. When
17 the volume multicast layer is called upon in a primary data mover to maintain remote
18 copies of a specified extent of a VolMCastSender instance, an instance of a
19 VolMCastCopy thread 374 is created and executed. The VolMCastCopy thread instance
20 accesses the delta sets from a primary save volume 375 to produce a write stream 376 of
21 blocks sent down to the RCP layer. At the secondary data mover, an instance of a
22 VolMCastReceiver thread 377 is instantiated and executed to receive a read stream 378
23 of blocks and write the copied delta sets into a secondary save volume 379. An instance

1 of an acknowledgement thread 380 returns an acknowledgement 381 of completion of
2 copying of a delta-set for an extent to the secondary file system. The acknowledgement
3 is sent down to the RCP layer of the secondary data mover. At the primary, the RCP
4 layer sends the acknowledgement 382 to an instance of an acknowledgement thread 383.

5 [000130] RCP is a session-layer protocol, for replication from one primary to
6 multiple secondary sites. Control is initiated by the primary, except when recovering
7 from aborted transfers. RCP uses TCP between the primary and secondary for control
8 and data. Network distribution is by an application-level multicast (ALM) using the RCP
9 as a forwarder. Port sharing with HTTP is used for crossing firewalls.

10 [000131] RCP may support other replication applications in addition to 1-to-N
11 IP-based replication for wide-area distribution of read-only data. These other
12 applications include 1-to-N volume mirroring, cluster file system commands, remote file
13 system replication, and distribution and replication of other commands that may be
14 recognized by the data movers.

15 [000132] The 1-to-N volume mirroring is a simplification of to 1-to-N IP-based
16 replication for wide-area distribution of read-only data, because the volume mirroring
17 need not synchronize a remote volume with any consistent version of the primary volume
18 until the remote volume needs to be accessed for recovery purposes.

19 [000133] Remote file system replication also uses RCP for broadcast or
20 forwarding an application command to a remote data mover to initiate a replication of a
21 file system managed by the remote data mover. In a similar fashion, RCP may broadcast
22 or forward other commands recognized by data movers, such as iSCSI or remote-control
23 type commands for archival storage. For example, RCP could broadcast or forward

1 remote control commands of the kind described in Dunham, U.S. Patent 6,353,878 issued
2 March 5, 2002 entitled "Remote Control of Backup Media in a Secondary Storage
3 Subsystem Through Access to a Primary Storage Subsystem," incorporated herein by
4 reference.

5 [000134] The RCP forwarder is composed of two RCP sessions, an outbound
6 session at the primary, and an inbound session at the secondary. The inbound RCP
7 session receives a group name and looks up the group in a routing table. If routes for the
8 group exist in the routing table, then an RCP forwarder is created at the secondary,
9 including a data path by pointer passing from an "in" session to an "out" session.

10 [000135] An RCP group may be configured to include application-level
11 multicast (ALM) topology. For example, ALM route configuration commands begin
12 with an identifier number for the network file server ("cel") that contains the forwarder
13 data mover, and an identifier number ("ser") for the forwarder data mover in the network
14 server. The configuration commands end with a "nexthop" specification of an immediate
15 destination data mover:

16
17 cell1-ser2: rcproute add group=g1 nexthop=cell2-ser2

18 cel2-ser2: rcproute add group=g1 nexthop=cell2-ser3

19 cel2-ser2: rcproute add group=g1 nexthop=cell2-ser4
20

21 [000136] In effect, the forwarder data mover adds the "nexthop" specification to
22 an entry for the RCP group in the routing table in the forwarder data mover. This entire
23 entry can be displayed by the following configuration command:

1

2 cel2-ser2: rcproute display

3

4 **[000137]** The entry is displayed, for example, as a list of the “nexthop”
5 destination data movers. The entry can be deleted by the following configuration
6 command:

7

8 cel2-ser2: rcproute delete

9

10 **[000138]** Each immediate destination data mover may itself be configured as a
11 forwarder in the RCP group. In this case, RCP commands and data will be forwarded
12 more than once, through a chain of forwarders. The set of possible RCP routes from a
13 primary or forwarder in effect becomes a tree or hierarchy of destinations.

14 **[000139]** The ALM commands may also include commands for creating
15 sessions and sending control blocks or data. For example, the following ALM command
16 creates a session and sends application data to all destinations in group “g1” from cell1-
17 ser2 from a file (named “filename”) using a test application (named “rcpfiletest”).

18

19 cell1-ser2: rcpfiletest data=filename group=g1

20

21 **[000140]** FIG. 14 shows the RCP collector service 390 at a primary site. The
22 programming for the RCP collector service includes an RCP session manager 391,
23 collector and worker threads 392, and a single-thread RCP daemon 393. The RCP
24 session manager 391 responds to requests from higher levels in the protocol stack, such

1 as a request from an application 394 to open an RCP pipe 395 between the application
2 394 and the RCP collector service 390. The application 394 may then send to the session
3 manager 391 requests to setup sessions with RCP groups. A session queue 396 stores the
4 state of each session, and a control block queue 397 keeps track of control blocks sent via
5 TCP/IP to the secondary data movers in the RCP groups. An RCP routing table 398
6 identifies the immediate destinations of each RCP group to which the TCP/IP messages
7 from the RCP collection service are to be sent, as well as any other destinations to which
8 the messages will be forwarded. For communication of the TCP/IP messages between
9 the RCP service and the network, TCP port :80 is opened in both directions (i.e., for input
10 and output). The single thread RCP daemon 393 is used for interfacing with this TCP
11 port :80.

12 **[000141]** FIG. 15 shows the RCP collector service 400 at a secondary site. The
13 RCP collector service at the secondary site is similar to the RCP collector service at the
14 primary site, in that it includes an RCP session manager 401, collector and worker
15 threads 402, a single thread RCP daemon 403 for access to/from TCP port :80, an RCP
16 session state queue 406, an RCP control block queue 407, and an RCP routing table 408.
17 The primary difference between the RCP collector service at the secondary site from the
18 RCP collector service at the primary site is in the collector and worker threads 402. At
19 the RCP secondary, the RCP commands and data are received from the TCP port :80
20 instead of from the application 404. The application 404 is the consumer of the RCP
21 data, instead of a source for RCP data. The RCP collector service 400 at the secondary
22 site may also serve as a forwarder for RCP commands, and therefore the RCP collector
23 service and worker threads 402 at the secondary site include a forwarder thread that does

1 not have a similar or complementary thread in the RCP collector service at the primary
2 site.

3 [000142] In operation, an application 404 can initialize the RCP collector
4 service so that the RCP collector service will call back the application upon receipt of
5 certain RCP commands from TCP port :80. For example, if a new connection command
6 is received from TCP port :80, then the RCP daemon 403 forwards the new connection
7 command to the RCP session manager. The RCP session manager 401 recognizes that
8 this connection command is associated with an application 404 at the secondary site,
9 opens an RCP pipe 405 to this application, and calls the application 404 indicating that
10 the RCP pipe 405 has been opened for the RCP session. (The volume multicast receiver
11 thread 377 of FIG. 13 is an example of such an application.) The application 404 returns
12 an acknowledgement. If the new connection is for a new RCP session, then the session
13 manager creates a new RCP session, and places state information for the new session on
14 the RCP session queue 406. RCP control blocks and data may be received for the session
15 from the TCP port :80. The data may be forwarded to the application, or to a file
16 specified by the application. RCP control blocks to be executed by the RCP collector
17 service 400 may be temporarily placed on the control block queue 407. RCP control
18 blocks or data intended for other secondary site may be forwarded to the intended
19 secondary sites.

20 [000143] FIG. 16 shows further details of the forwarding of RCP commands and
21 data by a data mover 430 identified as Cel2-Ser1. The data mover 430 is programmed
22 with a TCP/IP layer 431 for communication with the IP network 220, and an RCP layer
23 432 over the TCP/IP layer. For forwarding the RCP commands and data, the RCP layer

1 432 creates an inbound session 433 and an outbound session 434. The inbound session
2 433 receives RCP commands from the TCP/IP layer 431. The TCP/IP data stream is
3 retained in a data buffer 435. When an RCP command calls for the forwarding of RCP
4 commands or data to another data mover in a specified RCP group, the inbound session
5 433 performs a lookup for the group in a routing table 436.

6 [000144] In the example of FIG. 16, the routing table 436 includes a copy of all
7 of the routing information for each group of which the data mover 430 is a member. In
8 this case, for GROUP1, the primary data mover sends RCP commands to at least data
9 movers CEL2-SER1 and CEL9-SER1. CEL2-SER1 (i.e., the data mover 430) forwards
10 the RCP commands and RCP data to data movers CEL3-SER1 and CEL7-SER1. In
11 particular, the inbound session 433 creates an outbound session 434 and creates a TCP/IP
12 data path from the inbound session 433 to the outbound session 434 by passing pointers
13 to the data in the data buffer. The outbound session 434 invokes the RCP/IP layer 431 to
14 multicast the TCP data stream in the data buffer 435 over the IP network 220 to the data
15 movers CEL3-SER1 and CEL7-SER1.

16 [000145] The data mover CEL3-SER1 in succession forwards the RCP
17 commands to data movers CEL4-SER1 and CEL5-SER1. Normally, the data mover
18 CEL2-SER1 (430) does not need to know that the data mover CEL3-SER1 forwards the
19 RCP commands to data movers CEL4-SER1 and CEL5-SER1, but if the data mover
20 CEL2-SER1 (430) would fail to receive an acknowledgement from CEL3-SER1, then the
21 data mover CEL2-SER1 could minimize the impact of a failure of CEL3-SER1 by
22 forwarding the RCP commands to CEL4-SER1 and CEL5-SER1 until the failure of
23 CEL3-SER1 could be corrected.

1 **[000146]** FIG. 17 shows a flowchart of how the RCP collector service at the
2 secondary site processes an inbound RCP session command. In a first step 411, the RCP
3 collector service receives a session command. In step 412, if this session command is not
4 a command to be forwarded to other secondary sites, then execution branches to step 413
5 to execute the action of the command, and the processing of the session command is
6 finished.

7 **[000147]** In step 412, if the session command is a command to be forwarded to
8 other secondary sites, then execution continues from step 412 to step 414. In step 414,
9 the RCP collector service gets the RCP group name from the session command. Then, in
10 step 415, the RCP collector service looks up the group name in the RCP routing table
11 (408 in FIG. 15). If the group name is not found, then execution branches from step 416
12 to step 417. In step 417, the RCP collector service returns an error message to the sender
13 of the session command.

14 **[000148]** In step 416, if the group name is found in the RCP routing table, then
15 execution continues from step 416 to step 418. In step 418, the RCP collector service
16 forwards the action of the session command to each secondary in the group that is an
17 immediate destination of the forwarder (i.e., the data mover that is the secondary
18 presently processing the RCP session command). This is done by instantiating local
19 replication threads or creating outbound sessions for forwarding the action of the session
20 command to each secondary in the group that is an immediate destination of the
21 forwarder. After step 418, processing of the RCP session command is finished.

22 **[000149]** FIG. 18 shows an example of forwarding and local replication. In this
23 example, the IP network 220 connects a primary data mover 421 to a network file server

1 422 and a secondary data mover 423. The network file server 422 includes three data
2 movers 424, 425, and 426, and storage 427. The primary data mover manages network
3 access to a primary file system 428. The data mover 424 functions as a forwarder data
4 mover. The data mover 425 functions as a secondary data mover managing access from
5 the network to a secondary file system (copy A) 429. The data mover 426 functions as a
6 secondary data mover managing access from the network to a secondary file system
7 (copy B) 430. The data mover 423 manages network access to a secondary file system
8 (copy C) 431.

9 [000150] In operation, when the primary data mover 421 updates the primary
10 file system 428, it multicasts the modified logical blocks of the file system volume over
11 the IP network 220 to the forwarder data mover 424 and to the secondary data mover
12 423. The forwarder data mover 424 receives the modified blocks, and performs a local
13 replication of the blocks to cause the secondary data mover 425 to update the secondary
14 file system (copy A) 429 and the to cause the secondary data mover 426 to update the
15 secondary file system (copy B) 430.

16 [000151] To perform the local replication, the forwarder data mover 424 has its
17 volume multicast layer (323 in FIG. 9) save the modified blocks in a save volume 432 in
18 the storage 427, and then the forwarder data mover 424 sends replication commands to
19 the local secondary data movers 425 and 426. Each local secondary data mover 425, 426
20 has its playback module (328 in FIG. 9) replay the modifications from the save volume
21 432 into its respective secondary file system copy 429, 430.

22 [000152] FIG. 19 shows the sharing of the data mover's network TCP port :80
23 (440) between HTTP and RCP. This configuration is used in all data movers having the

1 RCP collector service; i.e., primary, secondary, or forwarder. The TCP data channel
2 from TCP port :80 (440) provides an in-order byte stream interface. IP packets 444 for
3 HTTP connections and IP packets 445 for RCP connections from the network 220 are
4 directed to the data mover's TCP port :80 (440). The TCP port :80 (440) is opened in
5 both directions (i.e., input and output). In the input direction, the data mover uses a level
6 5 (L5) filter 441 for demultiplexing the IP packets for the HTTP connections from the IP
7 packets for the RCP connections based on an initial segment of each TCP connection.
8 The L5 filter hands the TCP connection off to either a HTTP collector service 442 or an
9 RCP collector service 443. (The RCP collector service 443 is the collector service 390 in
10 the RCP primary of FIG. 14 or the RCP collector service 400 in an RCP secondary of
11 FIG. 15.) For example, if the initial segment of a TCP connection contains "HTTP/1.X",
12 then the L5 filter 441 directs the IP packets for the connection to the HTTP collector
13 service 442. If the initial segment of the TCP connection contains "RCP/1.0", then the IP
14 packets for the TCP connection are directed to the RCP collector service 443. (In an
15 alternative arrangement, the connection could be split as is done in a conventional stand-
16 alone IP switch.)

18 Data Recovery With Internet Protocol Replication With Or Without Full Resync

19 **[000153]** A remote replication system may protect against data loss resulting
20 from a disaster. For example, FIG. 20 shows the state of a remote replication system
21 before a disaster. The remote replication system copies data in an asynchronous fashion
22 from a primary data storage system at the primary site to a secondary data storage system
23 at a secondary site. The primary and secondary data storage systems could be file

1 servers, for example as shown in FIG. 1.

2 [000154] As shown in FIG. 20, the primary site includes a replication service
3 450 replicating changes made to a primary file system 451, and the secondary site
4 includes a playback service 452 writing the changes into a secondary file system 453.
5 The replication service 450 transfers the changes in “delta sets” sent over an IP pipe 454
6 to the playback service 452. The disaster occurs at the primary site during this transfer.
7 At this time, the playback service 452 is a few delta sets behind the replication service
8 450; for example, the playback service is playing back delta set (n-x) when the
9 replication service 452 is replicating delta set (n).

10 [000155] Once the disaster causes the primary site to go down, the secondary
11 site can be activated to service client read/write requests. When the primary site comes
12 back up, it can be re-synchronized to the secondary site, in many cases without making a
13 full copy of the file system. Then replication can be re-established as it was before the
14 disaster. This recovery process has been designed to minimize data loss and the down
15 time of both the primary and secondary site.

16 [000156] FIG. 21 is a flowchart of the preferred recovery process. Initially, in
17 step 461, data is replicated from the primary file system at the primary site by sending
18 delta sets to the secondary file system at the secondary site. Then in step 462, the
19 primary site becomes inoperative. In response, in step 463, the secondary site is activated
20 for read/write access to the secondary file system. This is done by a subroutine for
21 failover with a checkpoint and without sync, as further described below with respect to
22 FIG. 23. Eventually, in step 464, the primary site becomes operative. In response, in
23 step 465, the primary file system is synchronized to the state of the secondary file system,

1 as further described below with respect to FIG. 28. Then in step 466, read/write access to
2 the primary file system and replication of data from the primary file system to the
3 secondary file system is resumed in a failback operation, as further described below with
4 reference to FIG. 30.

5 [000157] FIG. 22 shows the state of the remote replication system of FIG. 20
6 during the failover operation. The (n-x) deltasets are played back into the secondary file
7 system 453, to put the secondary file system 453 into a state called a “restart point”. A
8 snapshot copy facility 456 creates a snapshot 455 of the “restart point” in order to
9 facilitate synchronization of the primary file system with the secondary file system once
10 the primary site becomes operative. The playback service 452 is stopped, and the
11 secondary file system 453 is remounted as read/write. The snapshot copy facility 456
12 keeps a record of changes made to the secondary file system 453 after the “restart point.”

13 [000158] FIG. 23 shows a flowchart of the failover operation. In step 471, if
14 possible, the primary file system is remounted as a read-only file system. Next, in step
15 472, the secondary site finishes playback of the (n-x) delta sets that it had received but
16 had not played back prior to the primary site becoming inoperative. In step 473, the
17 secondary site creates a snapshot copy of the “restart point” of the secondary file system.
18 This snapshot is stamped with the last delta set having been replayed into the secondary
19 file system prior to the creation of the snapshot. In other words, an identifier of this last
20 delta set is stored as an attribute of the snapshot. In general, whenever a snapshot is
21 made of a file system during playback of delta sets into the file system, the snapshot is
22 stamped with the last delta set having been replayed. In step 474, the secondary file
23 system is remounted as a read/write file system. In step 475, the snapshot process retains

1 and identifies all changes (delta) made to the secondary file system since the restarting
2 point.

3 [000159] Preferably the snapshot copy facility 456 retains and identifies the
4 changes at a logical volume level of data storage. For example, the present state of the
5 secondary file system is stored in a “clone volume,” and old versions of the logical blocks
6 that have been changed in the clone volume are saved in a “save volume”. In order to
7 conserve storage, the logical blocks of the save volume are dynamically allocated to the
8 old versions of the changed blocks as the changes are made to the clone volume. As
9 shown in FIG. 24, for each logical block that has been changed in the clone volume, a
10 block map 480 identifies the logical block address (S_i) of the old version of the block in
11 the save volume and the corresponding logical block address (B_i) of the changed block in
12 the clone volume.

13 [000160] FIG. 25 shows details of the preferred snapshot copy facility 456,
14 which provides multiple snapshots 483, 503 of a production file system 481. (Two
15 successive snapshots of the secondary file system are used, for example, in the resync
16 subroutine further described below with reference to FIG. 28.) The content of each
17 snapshot file system 483, 503 is the state of the production file system 481 at a particular
18 point in time when the snapshot was created. The snapshot copy facility 456 provides a
19 hierarchy of objects in a volume layer 490 supporting the file systems in a file system
20 layer 491. The production file system 481 is supported by read/write access to a file
21 system volume 482. Each snapshot file system 483, 503 provides read-only access to a
22 respective snapshot volume 484, 504.

1 **[000161]** Additional objects in the volume layer 490 of FIG. 25 permit the
2 content of each snapshot file system to be maintained during concurrent read/write access
3 to the production file system 481. The file system volume 482 is supported by a snapped
4 volume 485 having read access to a clone volume 487 and write access to a delta volume
5 486. The delta volume 486 has read/write access to the clone volume 487 and read/write
6 access to a save volume 488.

7 **[000162]** In the organization of FIG. 25, the actual data is stored in blocks in the
8 clone volume 487 and a respective save volume 488, 506 in storage for each snapshot.
9 The delta volume 486 also accesses information stored in a bit map 489 and the block
10 map 480. The bit map 489 indicates which blocks in the clone volume 487 have prior
11 versions in the save volume 488. In other words, for read-only access to the snapshot file
12 system, the bit map 489 indicates whether the delta volume should read each block from
13 the clone volume 487 or from the save volume 488. For example, the bit map is stored in
14 memory and it includes a bit for each block in the clone volume 487. The bit is clear to
15 indicate that there is no prior version of the block in the save volume 488, and the bit is
16 set to indicate that there is a prior version of the block in the save volume 488.

17 **[000163]** Consider, for example, a production file system 481 having blocks *a*,
18 *b*, *c*, *d*, *e*, *f*, *g*, and *h*. Suppose that when the snapshot file system 483 is created, the
19 blocks have values *a0*, *b0*, *c0*, *d0*, *e0*, *f0*, *g0*, and *h0*. Thereafter, read/write access to the
20 production file system 481 modifies the contents of blocks *a* and *b*, by writing new values
21 *a1* and *a2* into them. At this point, the following contents are seen in the clone volume
22 487 and in the save volume 488:

23

1 Clone Volume: *a1, b1, c0, d0, e0, f0, g0, h0*

3 Save Volume: *a0, b0*

5 [000164] From the contents of the clone volume 487 and the save volume 488, it
6 is possible to construct the contents of the snapshot file system 483. When reading a
7 block from the snapshot file system 483, the block is read from the save volume 488 if
8 found there, else it is read from the clone volume 487.

9 [000165] FIG. 25 further shows that a snapshot queue 500 maintains respective
10 objects supporting multiple snapshot file systems 483, 503 created at different respective
11 points in time from the production file system 481. In particular, the snapshot queue 500
12 includes a queue entry (J+K) at the tail 501 of the queue, and a queue entry (J) at the head
13 502 of the queue. In this example, the snapshot file system 483, the snapshot volume
14 484, the delta volume 486, the save volume 488, the bit map 489, and the block map 480
15 are all located in the queue entry at the tail 501 of the queue. The queue entry at the head
16 of the queue 502 includes similar objects; namely, the snapshot file system (J) 503, a
17 snapshot volume 504, a delta volume 505, a save volume 506, a bit map 507, and a block
18 map 508.

19 [000166] The snapshot copy facility 456 may respond to a request for another
20 snapshot of the production file system 481 by allocating the objects for a new queue
21 entry, and inserting the new queue entry at the tail of the queue, and linking it to the
22 snapped volume 485 and the clone volume 487. In this fashion, the save volumes 488,
23 506 in the snapshot queue 500 are maintained in a chronological order of the respective

1 points in time when the snapshot file systems were created. The save volume 506
2 supporting the oldest snapshot file system 503 resides at the head 502 of the queue, and
3 the save volume 488 supporting the youngest snapshot file system 483 resides at the tail
4 501 of the queue.

5 [000167] FIG. 26 shows a procedure performed by the snapshot copy facility for
6 writing a specified block (B_i) to the production file system. In step 511, if the snapshot
7 queue is not empty, execution continues to step 512. In step 512, the bit map at the tail of
8 the snapshot queue is accessed in order to test the bit for the specified block (B_i). Then in
9 step 513, if the bit is not set, execution branches to step 514. In step 514, the content of
10 the specified block (B_i) is copied from the clone volume to the next free block in the save
11 volume at the tail of the snapshot queue. Execution continues from step 514 to step 515.
12 In step 515, the save volume block address (S_i) of the free block is inserted into the entry
13 for the block (B_i) in the block map at the tail of the queue, and then the bit for the block
14 (B_i) is set in the bit map at the tail of the queue. After step 515, execution continues to
15 step 516. Execution also continues to step 516 from step 513 if the tested bit is found to
16 be set. Moreover, execution continues to step 516 from step 511 if the snapshot queue is
17 empty. In step 516, new data is written to the specified block (B_i) in the clone volume,
18 and then execution returns.

19 [000168] FIG. 27 shows a procedure performed by the snapshot copy facility for
20 reading a specified block (B_i) from a specified snapshot file system (N). In the first step
21 521, the bit map is accessed for the queue entry (N) to test the bit for the specified block
22 (B_i). Then in step 522, if the tested bit is set, execution continues to step 523. In step

1 523, the block map is accessed to get the save volume block address (S_i) for the specified
2 block (B_i). Then in step 524 the data is read from the block address (S_i) in the save
3 volume, and then execution returns.

4 **[000169]** If in step 522 the tested bit is not set, then execution branches to step
5 525. In step 525, if the specified snapshot (N) is not at the tail of the snapshot queue,
6 then execution continues to step 526 to perform a recursive subroutine call upon the
7 subroutine in FIG. 27 for read-only access to the snapshot ($N+1$). After step 526,
8 execution returns.

9 **[000170]** If in step 525 the snapshot (N) is at the tail of the snapshot queue, then
10 execution branches to step 527. In step 527, the data is read from the specified block (B_i)
11 in the clone volume, and execution returns.

12 **[000171]** FIG. 28 shows the subroutine to resync the primary file system with
13 the secondary file system, as used in the flowchart of FIG. 21. In a first step 531 in FIG.
14 28, the delta set identifier ($n-x$) for the restart point is read from the delta set attribute of
15 the restart point snapshot, and it is incremented by one to compute ($n-x+1$). The save
16 volume at the primary site is searched for the delta set ($n-x+1$). Execution continues from
17 step 531 to step 532 if the delta set $n-x+1$ is found in the save volume at the primary site.
18 In this case, the primary site should also have all of the delta sets from delta set $n-x+1$ to
19 delta set n , and the primary file system can be restored to the state of the restart point in
20 step 532 by an “undo” of the data blocks of these delta sets. The primary file system is
21 restored with “before images” of these data blocks in order to “undo” the changes. The
22 delta sets $n-x+1$ to n contain “after images” of these data blocks. The “before images” of
23 these data blocks are stored in the secondary site in the snapshot of the secondary file

1 system at the restart point. Therefore, to restore the primary file system to the state of the
2 restart point, the snapshot copy facility at the secondary site reads the “before images”
3 from the “restart point” snapshot, and the secondary site executes a remote copy of these
4 “before images” to the primary file system.

5 [000172] In step 532, the primary site restores the primary file system to the
6 state of the restarting point by obtaining a list of blocks from the save volume at the
7 primary site, including the blocks in delta set $n-x+1$ to delta set n . The primary site sends
8 this list to the snapshot copy facility at the secondary site. The secondary site retrieves
9 the data of these blocks from the snapshot at the restart point, and returns the data to the
10 primary site. The primary site receives these blocks of data and restores them to the
11 primary file system.

12 [000173] In step 533, the snapshot copy facility starts replication to the primary
13 file system, and creates an intermediate point snapshot (at state delta set $n-x+\delta$), so all
14 new changes made to the secondary file system since the intermediate point are kept and
15 sent to the primary file system. However, they are not yet played back into the primary
16 file system. At this point, the primary file system is mounted as a “raw” file system, so it
17 is not accessible to users.

18 [000174] In step 534, the changes made to the secondary file system from the
19 restart point to the intermediate point are copied from the secondary file system to the
20 primary file system. These changes are maintained by the snapshot copy facility at the
21 secondary site. For the snapshot copy facility as shown in FIG. 25, for example, assume
22 that the restart point is the snapshot J and the intermediate point is the snapshot J+K
23 where K is one. In this case, the changed blocks are identified in the bit map 507, and the

1 data of the changed data blocks are found by reading from the intermediate point
2 snapshot file system. To retrieve these changes, the snapshot copy facility scans the bit
3 map 507 to find any set bit indicating a change in the save volume, and if the bit is set,
4 the snapshot copy facility notes the block address (B_i). The snapshot copy facility then
5 reads the snapshot file system (J+K) 483 at this block address (B_i). The snapshot copy
6 facility continues this process until it has finished scanning the bit map 507.

7 **[000175]** In FIG. 28, after step 534, execution continues to step 535. In step
8 535, the primary file system is mounted as a read-only UxFS file system, and playback to
9 the primary file system is started. At this point, the resync is finished, and the primary
10 and secondary sites are ready for failback.

11 **[000176]** In step 531, if the delta set $n-x+1$ cannot be found in the save volume
12 from before failover (459 in FIG. 29) at the primary site (for example, because the save
13 volume was destroyed during the disaster), then execution branches from step 531 to step
14 536. In this case, the secondary file system is migrated to the primary file system, and a
15 warning message is logged and displayed to the system administrator to signal that the
16 full copy was required. In step 536, the secondary creates a first intermediate point
17 snapshot (at state delta set $n-x+\delta$), and copies the first intermediate point snapshot to the
18 primary file system. At this time, read/write access to the secondary file system is
19 permitted on a priority basis. Then in step 537, the secondary starts replication to the
20 primary file system, and creates a second intermediate point snapshot, so all new changes
21 made to the secondary file system since the second intermediate point are kept and sent to
22 the primary file system by the replication process. However, these new changes are not
23 yet played back into the primary file system. At this point, the primary file system is

1 mounted as a “raw” file system, so it is not accessible to users. Moreover, in step 537,
2 the differential of the first and second intermediate point snapshots (i.e., the changes
3 made to the secondary file system from the first intermediate point snapshot to the second
4 intermediate point snapshot) are copied to the primary file system. Then execution
5 continues to step 535, to mount the primary file system as a read-only UxFS file system,
6 and to start playback to the primary file system. Also in step 535, the restart snapshot
7 and the intermediate snapshot can be deleted.

8 [000177] The migration method of steps 536 to 537 is used in order to reduce
9 the likelihood that read/write access to the secondary file system and replication to the
10 primary file system would be interrupted by the primary and secondary save volumes
11 (274, 275 in FIG. 5) becoming full. Each save volume functions as a wrap-around buffer.
12 The primary save volume will store updates received in step 537 from the secondary site
13 until playback to the primary file system is started in step 535. If the primary save
14 volume would become full in step 537, a TCP flow control signal is sent from the
15 primary site to the secondary site to suspend transmission of further updates. When
16 transmission of further updates is suspended, the further updates can be stored in the
17 secondary save volume (275 in FIG. 5) until the secondary save volume would become
18 full. If the secondary save volume would become full, then read/write access to the
19 secondary file system and replication to the primary file system would be suspended until
20 playback to the primary file system is started in step 535 and updates are removed from
21 the primary and secondary save volumes.

22 [000178] FIG. 29 shows the state of the remote replication system of FIG. 20
23 during the resync procedure. At the secondary site, the snapshot copy facility 456 has

1 created the intermediate point snapshot 457. A list of blocks in the delta sets n to $n-x+1$
2 is obtained from the save volume 459 at the primary site and sent to the snapshot copy
3 facility 456 at the secondary site. The snapshot copy facility 456 returns “before images”
4 of the requested blocks over an IP pipe 458 to the primary file system 451 to restore the
5 primary file system to the state of the restart point snapshot, and then sends the changes
6 from the restart point to the intermediate point. The replication service 452 at the
7 secondary site is replicating changes to the secondary file system since the intermediate
8 point, and is sending these changes over the IP pipe 460 to the playback service 450.
9 Transmission of these changes since the intermediate point over the IP pipe 460 is
10 concurrent with transmission of changes prior to the intermediate point over the IP pipe
11 458. Once the primary file system 451 has been synchronized to the intermediate point
12 457, the playback service 450 at the primary site is activated to play these changes into
13 the primary file system 451. If the primary site should crash during the copy process of
14 steps 534 or 536, then the copy process can be restarted automatically.

15 [000179] FIG. 30 is a flowchart of the subroutine to failback to the primary file
16 system. In a first step 541, the primary file system is made almost identical to the
17 secondary file system by the replication process. For example, when the primary file
18 system is within one delta chunk of the secondary file system, execution continues to step
19 542. In step 542, the secondary file system is re-mounted as read-only, and the last delta
20 chunk is copied from the secondary file system to the primary file system in order to
21 synchronize the primary file system from the secondary file system. Once synchronized,
22 the primary and secondary are aborted, in order to terminate the processes of replication
23 from the secondary, playback to the primary, and copying from the secondary to the

primary. At this point, the primary file system contains all the data that the secondary file system had before the disaster and the additional data that the secondary file system had accumulated after the disaster. Then in step 543, there is a resumption of the replication of data from the primary file system and playback to the secondary file system. Then in step 544, the primary file system is remounted as read/write. Therefore, the recovery process permits replication to be restarted as it was before the disaster. Some data might have been lost at failover time since replication is asynchronous. This is dependent on the number “x” of delta sets that the secondary file system was behind the primary file system when the disaster occurred at the primary. The less the number “x” of delta sets the closer the final recovered file system would be to the original file system.

[000180] A number of commands have been devised for conducting the above recovery procedure when network file servers (such as shown in FIGS. 1 and 5) are used at the primary and secondary sites. These commands include a `$fs_replicate-failover` command, a `$fs_replicate -resync` command, a `$fs_replicate-status` command, `$fs_replicate-failback` command, and a `$fs_copy -start` command.

The `$fs_replicate-failover` command has the following format:

```
$fs_replicate -failover <pfs>:cel=<server> <sfs> [-sync] [-now]
```

where `<pfs>` is the primary file system, `<server>` is the server of the primary file system, `<sfs>` is the secondary file system, `[-sync]` is an option that is set when the user wants the failover to happen when both the primary file system and the secondary file system are

1 “in sync” (i.e., identical state), and [-now] is an option to specify immediate failover.
2 The \$fs_replicate-failover command is executed on the control station for the secondary
3 file system.

4 **[000181]** When the [-sync] option is not set, this command will internally stop
5 the replication and playback of <sfs>, making sure that all of the delta sets available on
6 secondary site are re-played. Then it will try to stop the replication on <pfs> as well,
7 unless the [-restartpoint] option is set. Then <sfs> will be remounted “rw”. If possible,
8 <pfs> is remounted “ro”. No [-sync] option should be specified if the primary control
9 station or server is down. In this case, after failover, some data could be lost.

10 **[000182]** When the [now] option is set, the playback of <sfs> is immediately
11 stopped, without replaying any of the delta sets available on the secondary site.

12 **[000183]** When the [-sync] option is specified, the primary control station and
13 the server must be up, else the command will fail. The primary file system <pfs> is re-
14 mounted as read-only and a last delta set is created. Once the last delta set is played back
15 to the secondary, the replication process is stopped on the <sfs>/<pfs>. The secondary
16 file system <sfs> now contains all the data that the <pfs> had before the failover. The
17 secondary file system <sfs> is now mounted “rw”.

18 **[000184]** FIG. 31 shows a flowchart of execution of the \$fs_replicate-failover
19 command. In a first step 551, if the sync option is not set, then execution branches to
20 step 552 to perform a failover without sync procedure, as described below with reference
21 to FIG. 32. After step 552, execution is finished. In step 551, if the sync option is set,
22 then execution continues to step 553. In step 553, if the primary site is not operational,
23 execution returns reporting a command failure. Otherwise, if the primary site is

operational, execution continues to step 554, to perform a failover with sync procedure, as described below with reference to FIG. 33.

[000185] FIG. 32 shows a flowchart of the failover without sync procedure. In step 561, if possible, the primary file system is remounted as a read-only file system. Then in step 562, the secondary site finishes playback of the delta sets that it had received but had not played back prior to receipt of the failover command. In step 563, the secondary site creates a snapshot copy of the secondary file system. This snapshot is called the restarting point of the file system. In step 564, the secondary file system is remounted as a read/write file system, and execution of the failover command is finished.

[000186] FIG. 33 shows a flowchart of the failover with sync procedure. In a first step 571, the secondary file system is made almost identical to the primary file system by the replication process. For example, once the secondary file system is within one delta set chunk of the primary file system, execution continues to step 572. In step 572, the primary file system is remounted as read-only, the last delta chunk is copied from the primary file system to the secondary file system in order to synchronize the secondary file system to the primary file system, and then the primary and secondary are aborted to terminate the processes of replication from the primary, playback to the secondary, and copying from the primary to the secondary. Then in step 573, the secondary site creates a snapshot copy of the secondary file system. This snapshot copy is called the restarting point of the file system. Finally, in step 574, the secondary file system is remounted as read/write.

[000187] The `$fs_replicate-resync` command has the following format:

1 \$fs_replicate -resync sfs pfs:cel=server1

2
3 where sfs is the secondary file system (this is now the source file system), pfs is the
4 primary file system (this is the file system to be restored, a raw file system mounted on a
5 server which has its IP interfaces configured), server1 is the site where a disaster
6 occurred, and <ckpt_name> is the restart point name.

7
8 **[000188]** The \$fs_replicate-status command has the following format:

9
10 \$fs_replicate -status <pfs> <sfs>:cel=<remote_server>

11
12 where <pfs> is the primary file system, <sfs> is the secondary file system, and
13 <remote_server> is the server of the secondary site. This command can be used on the
14 control station for the primary site, and also on the control station for the secondary site.
15 Typically this command is used on the primary site when the user wants to check the
16 status of the replication process, to determine when it is reasonable to failback. It is
17 considered reasonable to failback when the difference between the primary file system
18 and the secondary file system is very small. To assist the user, the primary site responds
19 to this command by providing an indication of the number of blocks that are different
20 between the primary file system and the secondary file system.

21 **[000189]** Following is an example of the information provided when the
22 \$fs_replicate-status command is executed:

```

1      [nasadmin@cel_sfs nasadmin]$ fs_replicate -info sfs10
2
3      id                      = 59
4
5      name                    = sfs10
6
7      type                    = replication
8
9      current_delta_set       = 16
10
11     current_block_number     = 1088
12
13     current_transfer_rate    = 17.0667 blocks/second
14
15     avg_transfer_rate        = 25.7444 blocks/second
16
17     id                      = 31
18
19     name                    = pfs10:cel_ip2
20
21     type                    = playback
22
23     current_delta_set       = 16
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99

```

outstanding delta sets:

Delta	Source_create_time	Blocks
15	11/26 20:11:56	1

Delta	Create Time	Dur	Blocks	Playback	Dur	Blocks	Dsin	Group
15	11/26 20:11:56	0	1					
14	11/26 20:01:56	0	1					

1	13	11/26 19:51:56 0	1	11/26 21:08:48 0	1	2
2	12	11/26 19:41:56 0	1			
3	11	11/26 19:31:56 0	1	11/26 20:48:48 0	1	2
4	10	11/26 19:21:56 0	1			
5	9	11/26 19:11:56 0	1	11/26 20:28:28 0	1	2
6	8	11/26 18:52:40 0	1	11/26 20:08:09 0	1	1
7	7	11/26 18:42:40 0	1			
8	6	11/26 18:32:40 0	1	11/26 19:48:09 0	1	2
9	5	11/26 18:22:39 0	1			
10	4	11/26 18:12:39 0	1	11/26 19:28:08 0	1	2
11	3	11/26 18:02:39 0	4	11/26 19:07:23 0	4	1

12

13 **[000190]** The \$fs_replicate-failback command has the following format:

14

15 \$fs_replicate -failback <pfs> <sfs>:cel=<remote_server>

16

17 where <pfs> is the primary file system (a UxFS file system mounted as “ro”), and <sfs>
 18 is a secondary file system. This command is used when the user wants to failback to the
 19 primary site (after re-build phase). No data is lost and the replication will continue but it
 20 will be reversed (the primary file system, which runs playback, will then run replication
 21 and vice-versa for the secondary file system). This command can also be used to change
 22 the direction of replication, in the absence of any failure or disruption.

23 **[000191]** The \$fs_replicate-failback command is executed on the site that the
 24 replication service is running (i.e., the site having the file system that is mounted as

1 read/write). Typically, the `$fs_replicate-failback` command is executed on the primary
2 site after it has been rebuilt. Then the primary file system is a raw file system that is
3 restored from the secondary file system. The secondary file system is re-mounted read
4 write. Replication is stopped on the secondary and playback is stopped on the primary.
5 After the primary file system and the secondary file systems are in sync, the primary file
6 system is converted to a UxFS file system and re-mounted as read/write. Replication is
7 turned back on and the flow of data is now from the primary file system to the secondary
8 file system. Playback is started on the secondary file system. This command can be used
9 to “swap” primary and secondary without having to do a full copy and without having to
10 restart replication. On error, it should be possible to re-issue the `fs_replicate -failback`
11 command and proceed if the error condition has been cleared.

12 **[000192]** The `$fs_copy-start` command has the following format:

13
14 `$fs_copy -start <fs_ckpt2> <sfs>:cel=<remote_server> -fromfs <fs_ckpt1>`
15

16 where `<fs_ckpt1>` is the primary file system ckpt1, `<fs_ckpt2>` is the primary file system
17 ckpt2, `<sfs>` is the secondary file system, and `<remote_server>` is the server of the
18 secondary site. This command is used internally by the `fs_replicate-resync` command. It
19 could also be used for other purposes, for example, as described below with reference to
20 FIGS. 36-38. This command copies the delta from `fs_ckpt1` to `fs_ckpt2` over to the
21 secondary file system, knowing that the state of the secondary file system was equal to
22 `fs_ckpt1`.

1 **[000193]** The following is an example of how the commands can be used to
2 recover from a disaster.

3 1. After the disaster, decide the option to use for failover.

4 2. \$fs_replicate -failover pfs:cel=cel_pfs sfs

5 3. Verify that sfs can accept read/write operation.

6 4. Initialize replication from sfs to pfs:

7 On PFS control station:

8 server_mount server_2 -o ro pfs /pfs

9 On SFS control station:

10 \$fs_replicate -resync pfs:cel=cel_pfs sfs

11 5. Verify that replication from sfs to pfs is running without a problem.

12 6. \$fs_replicate -failback sfs pfs:cel=cel_pfs

13 7. Verify that replication from pfs to sfs is running without a problem.

14 15 Replication of Snapshots Using IP File System Copy Differential

16 **[000194]** As described above with reference to steps 534 and 537 of FIG. 28, a
17 file system copy differential was used to replicate an intermediate point snapshot during
18 recovery of a file system from a disaster. A general-purpose file system copy differential
19 command (\$fs_copy-start) was also introduced.

20 **[000195]** FIG. 34 shows a flowchart of a procedure for the \$fs_copy -start
21 command. Here L is the index of the snapshot fs_ckpt1, and M is the index of the
22 snapshot fs_ckpt2. In a first step 581, M is compared to L, and execution branches to

1 step 582 to return an error if M is not greater than L. Otherwise, execution continues to
2 step 583.

3 [000196] The following steps execute the \$fs_copy –start command by scanning
4 one or more of the bit maps of the snapshot copy facility of FIG. 25. Each bit map
5 indicates whether or not each block of a respective snapshot changed from the snapshot
6 to the next snapshot. Therefore, by examining the bit maps for the snapshots L to M-1, it
7 is determined whether or not each block has changed from snapshot L to M. If so, the
8 new value of the block is read from the snapshot M. The bit maps are scanned in block
9 order so that the new blocks are read from the snapshot M of the primary (i.e., the source)
10 file system and transmitted to the secondary (i.e., the destination) file system in the order
11 of their block indices (B_i).

12 [000197] In step 583, the block index (B_i) is set to zero. In step 584, a snapshot
13 index (I) is set to L. In step 585, the bit map for the snapshot (I) is indexed with the
14 block index (B_i) to determine whether or not the block was changed between snapshot (I)
15 and snapshot (I+1). If in step 585 the bit for the block (B_i) is not set in the bit map for the
16 snapshot (I), then no such change occurred, and execution continues to step 586. In step
17 586, if the snapshot index I is not equal to M-1, then execution continues step 587. In
18 step 587, the snapshot index I is incremented by one. After step 587, execution loops
19 back to step 585. In step 585, if the bit for the block (B_i) is set in the bit map for the
20 snapshot (I), then execution branches to step 589. In step 589, the snapshot (M) is read to
21 get the new data for the block (B_i). Execution continues to step 590 of FIG. 35. In step
22 590, the block index (B_i) and the new data for the block (B_i) from the snapshot (M) are

1 returned for copying into the secondary file system. Then in step 591, if the block index
2 (B_i) is at the end of the production volume, then execution is finished. If not, execution
3 loops back to step 588 of FIG. 34. In step 588, the block index (B_i) is incremented by
4 one, and execution loops back to step 584, to continue scanning for the next value of the
5 block index.

6 [000198] In step 586 of FIG. 34, if I is equal to $M-1$, then execution continues to
7 step 591 of FIG. 35.

8 [000199] It should be understood that the flowchart of FIGS. 34-35 represents a
9 program executed by at least one processor in a data storage system such as a network
10 file server. The processor, for example, is a data mover computer (e.g., 232 in FIG. 1).
11 The program, for example, is initially contained in a program storage device such as a
12 floppy disk (e.g., 238 in FIG. 1) and down-loaded into storage of the data mover
13 computer.

14 [000200] The program in the flowchart of FIGS. 34-35 has an inner loop
15 including steps 585, 586, 587 that indexes the snapshots L to snapshot $M-1$. This
16 sequence includes the snapshot L and the snapshots that are both younger than the
17 snapshot L and older than the snapshot M . The program in the flowchart of FIGS. 34-35
18 has an outer loop including steps 584, 585, 586, 591, and 588 that indexes the blocks.
19 When a bit in the indexed bit map is found to be set in step 585, the inner loop is exited to
20 return the block index (B_i) and the data in the snapshot M for block (B_i).

21 [000201] The snapshot copy differential has been described above for
22 facilitating recovery of a file system after a disaster. The snapshot copy differential can
23 also be used for wide-area distribution of updates on an as-needed basis. This reduces

1 network traffic for the case where a client has an old local version of a file system and
2 needs a new version of the file system. A new local version of the file system can be
3 constructed by copying the appropriate changes into the old local version of the file
4 system.

5 [000202] FIG. 36, for example, shows a block diagram of a data network in
6 which snapshot deltas are transmitted over a wide-area network 626 from a network file
7 server 627 to a local file server 624 in order to update a local file system 625 as needed.
8 The local file server 624 services local clients 621, 622, and 623. When a client needs a
9 more recent version of the file system, and the local file system 625 is not the most recent
10 version, then the local file server may request a specified version from the network file
11 server 627, or the client may request the most recent version available.

12 [000203] The network file server 627 has a snapshot copy facility 628 storing
13 multiple snapshots 629, 630. If the local file system 625 in the local file server 624 is one
14 of the multiple snapshots, then the network file server 627 may respond to a request from
15 the local file server 624 by obtaining from the snapshot copy facility 628 a snapshot copy
16 differential that would contain all of the updates necessary to convert the local file system
17 624 to a more recent snapshot of the file system. In the usual case, the local file server
18 624 would request all of the updates necessary to convert the local file system 625 to the
19 most recent snapshot copy. However, it is also possible for the local file server 624 to
20 request the updates for a specified version that would not necessarily be the most recent
21 snapshot copy.

22 [000204] FIGS. 37 and 38 show a flowchart of a procedure for the replication of
23 the most recent snapshot in the system of FIG. 36 using the snapshot copy differential. In

1 a first step 641, the client requests access to the file in the local file system in the local
2 file server. In step 642, the local file server accesses attributes of the local file system
3 and finds that it is version (Q) of a local updatable file system, last updated at time (Tu)
4 from a network file server having a network identifier (NETID). The local file server
5 uses the time of last update (Tu) to determine that it is time to check for a more recent
6 version, and sends an update request to the network file server (NETID). For example,
7 the file system has an attribute specifying a minimum update interval (Tmin), and it is not
8 time to request an update unless the minimum update interval has elapsed since the time
9 of the last update. The request specifies the version (Q) already in the local file server.
10 In response, the network file server accesses the snapshot copy facility to find the oldest
11 version (J) and the youngest version (J+K) stored in the network file server. In step 645,
12 Q is compared to J+K. If $Q=J+K$, then execution branches to step 646 of FIG. 38. In
13 step 646, the network file server returns a message that no more recent version is
14 available. The local file server resets the time of last update (Tu) to the current time, and
15 accesses the local version (Q) for the client. Upon completion of step 646, the snapshot
16 replication process is finished.

17 [000205] If in step 645 Q is not equal to J+K, then execution continues to step
18 647 of FIG. 38. In step 647, Q is compared to K. If Q is less than K, then execution
19 branches to step 648. In this case, the version (Q) is not a snapshot in the snapshot copy
20 facility because the version (Q) is too old. In step 648, the network file server copies the
21 youngest snapshot version (J+K) to the local file server. The local file server replaces the
22 local version (Q) with the new version (J+K), resets the time of last update (Tu) to the

1 current time, and accesses the new local version (J+K) for the client. Upon completion of
2 step 648, the snapshot replication process is finished.

3 [000206] In step 647, if Q is not less than K, then execution continues to step
4 649. In step 649, the network file server does a file system copy snapshot delta <Q> to
5 <J+K> of blocks into the local version (Q) to convert it into the youngest snapshot
6 version (J+K). The local file server resets the time of last update (Tu) to the current time,
7 and accesses this local snapshot version (J+K) for the client. Upon completion of step
8 649, the snapshot replication process is finished.

9 [000207] In a preferred snapshot copy facility, as described below with
10 reference to FIGS. 41 to 46, there is kept a meta bit map for each snapshot copy for
11 indicating blocks of the production file system that are not used in the snapshot copy.
12 Further details of such a snapshot copy facility are described in Philippe Armangau, et
13 al., "Data Storage System Having Meta Bit Maps for Indicating Whether Data Blocks are
14 Invalid in Snapshot Copies," U.S. Patent Application Ser. 10/213,241 filed Aug. 6, 2002,
15 incorporated herein by reference. The snapshot copy facility maintains the meta bit maps
16 in order to store the "before image" of a block in the save volume at the tail of the
17 snapshot queue only when the block is being written to and the "before image" is needed
18 for responding to any request for reading a snapshot copy. This reduces the number of
19 blocks that are stored in the save volumes. However, in this case, the bit map for each
20 snapshot (L) indicates whether or not a block has been stored in the save volume for the
21 snapshot (L), and no longer will indicate all of the blocks that have been changed after
22 snapshot (L) and before snapshot (L+1). In particular, if a block was not in use for

1 snapshot (L), and was written to after snapshot (L) and before snapshot (L+1), then the
2 “before image” of the block will not be written into the save volume for snapshot (L).

3 [000208] When it is known that a block is not used in the snapshot copy (M),
4 then there is no need for the snapshot copy facility to return the block when responding to
5 a request for the snapshot delta of snapshot <L> to snapshot <M>. Therefore, for the
6 preferred snapshot copy facility, it is desirable to modify the procedure of FIG. 34 in
7 order use the information in the meta bit map for the snapshot <M>. In this case, the
8 procedure of FIG. 34 should also be modified to account for the fact that the save
9 volumes no longer store the “before images” for all of the blocks that may have changed
10 between the successive snapshot copies.

11 [000209] FIG. 39 shows how the flowchart of FIG. 34 can be modified for use
12 with the preferred snapshot copy facility of FIGS. 41 to 46. Steps 651, 652, and 653 of
13 FIG. 39 are similar to steps 581, 582, and 583 of FIG. 34. Step 654 of FIG. 39 is similar
14 to step 588 of FIG. 34, and step 658 of FIG. 39 is similar to step 589 of FIG. 34.

15 [000210] In FIG. 39, after step 653 or step 654, execution continues to step 655.
16 In step 655, if the block (B_i) is not in use in the snapshot (M), then execution branches to
17 step 591 of FIG. 35. Therefore, a block not in use in the snapshot (M) is not returned in
18 response to the command to copy the snapshot delta <L> to <M>. If the block (B_i) is in
19 use in the snapshot (M), then execution continues from step 655 to step 656. In step 656,
20 if the block (B_i) is in any of the save volumes (L) to (M-1), then execution continues to
21 step 658 to read the block (B_i) from the snapshot (M) in order to return this version of the
22 block (B_i) in response to the command to copy the snapshot delta <L> to <M>. In step
23 656, if the block (B_i) is in any of the save volumes (L) to (M-1), then execution branches

1 to step 657. In step 657, if the block (B_i) is in use in all of the snapshots (L) to (M-1),
2 then execution branches to step 591 of FIG. 35. In this case, the block (B_i) did not
3 change from snapshot (L) to snapshot (M), because if it did, a “before image” would have
4 been stored in one of the save volumes (L) to (M-1). In step 657 of FIG. 39, if it is not
5 true that the block (B_i) is in use in all of the snapshots (i.e., it is true that the block (B_i) is
6 not in use in at least one of the snapshots (L) to (M-1)), then execution continues from
7 step 657 to step 658 to read the block (B_i) from the snapshot (M) in order to return this
8 version of the block (B_i) in response to the command to copy the snapshot delta <L> to
9 <M>. In this case, it is possible that the block changed from snapshot (L) to snapshot
10 (M) despite the fact that the block (B_i) is not in any of the save volumes (L) to (M-1).

11 **[000211]** FIG. 40 shows a preferred implementation of the procedure of FIG. 39.
12 Steps 661 to 665 of FIG. 40 are similar to steps 651 to step 655 of FIG. 39. In step 665,
13 the meta bit map for snapshot (M) has a value for the block (B_i) indicating whether or not
14 the block (B_i) is in use for the snapshot (M). In particular, a value of 1 indicates that the
15 block (B_i) is in use for the snapshot (M). Steps 666, 669, and 670 of FIG. 40 are similar
16 to step 584, 586 and 587 of FIG. 34. From step 666 or step 670, execution continues to
17 step 667.

18 **[000212]** In step 667, if the bit map for snapshot (I) has a value of 1 for the
19 block (B_i), then execution continues to step 671 to read the snapshot (M) to get data for
20 the block (B_i), in order to return the data in response to the command to copy the
21 snapshot delta <L> to <M>. In this case, the save volume for block (I) includes a “before
22 image” for the block (B_i). Otherwise, if the bit map for snapshot (I) does not have a

1 value of 1 for the block (B_i), execution branches from step 667 to step 668. In step 668,
2 if the meta bit map for the snapshot (I) does not have a value of 1 for the block (B_i),
3 execution continues to step 671 to read the snapshot (M) to get data for the block (B_i), in
4 order to return the data in response to the command to copy the snapshot delta <L> to
5 <M>. In this case, the block (B_i) is not in use in the snapshot (I). Otherwise, if the meta
6 bit map for the snapshot (I) has a value of 1 for the block (B_i), execution continues to step
7 669.

9 Maintenance of Meta Bit Maps in the Snapshot Copy Facility

10 **[000213]** In the above description of the snapshot copy process, and in particular
11 with respect to FIG. 25, it was assumed that the original contents of a block of the
12 production file system must be saved to the most recent save volume before the contents
13 of the block are modified by a write access to the production file system. In practice,
14 however, the original contents are often invalid, and therefore need not be saved. For
15 example, many applications start with an empty file, and the file increases in size as data
16 is written to the file. In some of these applications, the file rarely decreases in size.
17 However, storage for the file may be released when the file is deleted from the file server,
18 for example, when the file is transferred to archival storage. In some applications, the
19 extent of a file may be dynamically decreased concurrent with read/write access to the
20 file.

21 **[000214]** There are significant advantages to identifying when read/write access
22 to the production file system is about to modify the contents of an invalid data block. If
23 this can be done in an efficient manner, then there can be a decrease in the access time for

1 write access to the production file system. A write operation to an invalid block can be
2 executed immediately, without the delay of saving the original contents of the data block
3 to the most recent save volume at the tail of the snapshot queue. Moreover, there is a
4 saving of storage because less storage is used for the save volumes. There is also a
5 decrease in memory requirements and an increase in performance for the operations upon
6 the snapshot file systems, because smaller bit and block hash indices can be used, and the
7 reduced amount of storage for the snapshots can be more rapidly restored to the
8 production file system, or deallocated for re-use when snapshots are deleted.

9 [000215] An efficient way of identifying when read/write access to the
10 production file system is about to modify the contents of an invalid data block is to use a
11 meta bit map having a bit for indicating whether or not each allocated block of storage in
12 the production file system is valid or not. For example, whenever storage is allocated to
13 the production file system by the initial allocation or the extension of a clone volume, a
14 corresponding meta bit map is allocated or extended, and the bits in the meta bit map
15 corresponding to the newly allocated storage are initially reset.

16 [000216] FIG. 41 shows a procedure for writing a specified block (Bi) to the
17 production file system when there is a meta bit map for indicating invalid data blocks in
18 the production file system. In a first step 681, a queue pointer is set to point to the queue
19 entry at the tail of the snapshot queue. Next, in step 682, the bit map in this snapshot
20 queue entry is accessed to test the bit for the specified block (Bi). Next, in step 683, if
21 the tested bit is found to be set, then execution continues to step 684. In step 684, new
22 data is written to the specified block (Bi) in the clone volume, and then execution returns.

1 **[000217]** In step 683, if the bit in the bit map is not set, then execution branches
2 to step 685. In step 685, the meta bit map in the snapshot queue entry is accessed to test
3 the bit for the specified block (Bi). Then, in step 686, execution continues to step 687 if
4 this bit is found to be set. In step 687, the content of the block (Bi) is copied from the
5 clone volume to the next free block in the save volume at the tail of the snapshot queue.
6 In step 688, an entry for the block (Bi) is inserted into the block map at the tail of the
7 snapshot queue, and then the bit for the block (Bi) is set in the bit map at the tail of the
8 snapshot queue. Execution continues from step 688 to step 684, to write new data to the
9 block (Bi) in the clone volume.

10 **[000218]** In step 686, if the tested bit is found not to be set, then execution
11 branches to step 689. If the queue pointer is pointing to the head of the queue, then
12 execution branches to step 684, to write new data to the block (Bi) in the clone volume.
13 Otherwise, if the queue pointer is not pointing to the head of the snapshot queue, then
14 execution continues to step 690 to advance the queue pointer to the next snapshot queue
15 entry toward the head of the snapshot queue. After step 690, execution loops back to step
16 682.

17 **[000219]** FIG. 42 shows an organization of the snapshots in the network file
18 server when a respective meta bit map 692, 693 is maintained for each snapshot in
19 addition to the meta bit map 691 for the production volume. It is desired to maintain a
20 respective meta bit map for each snapshot so that whenever the production file system is
21 restored with a snapshot file system, the meta bit map for the production file system can
22 be restored with the meta bit map for each snapshot. For example, a meta bit map of the
23 production file system is included in the production file system, so that whenever a

1 snapshot copy of the production file system is created, a snapshot copy of the meta bit
2 map is also created. Consequently, when the production file system is restored with a
3 snapshot, the meta bit map of the production volume is replaced with the meta bit map of
4 the snapshot.

5 [000220] As shown in FIG. 42, a meta bit map 691 is linked to the production
6 file system 481 for indicating invalid blocks in the production file system. Each entry in
7 the snapshot queue 500 includes a respective meta bit map linked to the snapshot file
8 system in the entry. For example, the queue entry (J+K) at the tail 501 of the queue has a
9 meta bit map 692 linked to the snapshot file system 483, and the queue entry (J) at the
10 head 502 of the queue includes a meta bit map 693 linked to the delta volume 505.

11 [000221] To reduce the memory and storage requirements for maintaining the
12 bit maps 694, 696 and block maps 695, 697, the each bit map is organized as a set of
13 pages indexed by a page table, and the each block map is organized as a set of hash lists
14 indexed by a hash table. The bit maps and block maps 694, 695 at the queue entry (J+K)
15 at the tail of the queue are initially created in a random access memory cache and written
16 back to storage of the save volume 488 when a next snapshot of the production file
17 system 481 is created. Thereafter the bit maps and block maps can be deallocated from
18 the random access memory in order to free up cache memory, and later staged from
19 storage to the cache memory when needed for read access to their respective snapshots.

20 [000222] FIG. 43 shows that the bit map 694 is organized as a page table 701
21 indexed by a set of most significant bits (MSB) of the block index (Bi), and pages 702
22 and 703 of the bit map linked to respective entries of the page table. The page table 701
23 includes a set of entries, each of which is either zero, indicating that the entire page is

1 zero and therefore the page is not stored in storage or random access memory, or is a
2 pointer to a page of the bit map.

3 [000223] FIG. 44 shows that the block map 695 is organized as a hash table 708
4 indexed by a hashing of the block index (B_i), and a set of hash lists 709 linked to
5 respective entries of the hash table 708. Each non-zero entry in the hash table 708 points
6 to a respective one of the hash lists 709. Each entry in each hash list includes a block
7 address (B_i) to a block in the clone volume, a corresponding block address (S_i) of the
8 block in the save volume, and a value that is either zero indicating the end of the has list,
9 or a pointer to the next entry in the list.

10 [000224] It is not necessary to retain the bit map 694 for the snapshot because
11 the block map 695 can be accessed to provide the same information that is contained in
12 the bit map 694. In particular, the block map 695 can be accessed simply to determine
13 whether or not a specified block (B_i) is found in the block map, and therefore can be
14 found in the save volume for the corresponding snapshot. However, the bit map 694 can
15 be accessed more quickly to provide this information. Therefore, for applications such as
16 backup where the snapshots would be accessed very infrequently, it may be desirable to
17 conserve storage by discarding the bit map for the snapshot at the tail of the queue once
18 the next snapshot is created. Otherwise, for an application such as described above with
19 respect to FIGS. 36-38 where the snapshots would be accessed frequently, the bit map for
20 the snapshot at the tail of the snapshot queue may be written to storage and maintained in
21 a random access cache memory as needed.

22 [000225] FIG. 45 shows that the meta bit map 691 has a respective bit
23 corresponding to each block in the clone volume, and in this example, each bit in the

meta bit map corresponds to one and only one block in the clone volume. The meta bit map 691 includes a series of words, each with a multiple of M bits. In this example, a bit having a value of zero indicates a corresponding block that is invalid, and a bit having a value of one indicates a corresponding block that is valid.

[000226] The meta bit map, however, may have a granularity greater than one block per bit. For example, each bit in the meta bit map could indicate a range of block addresses, which may include at least some valid data. The benefit to the increased granularity is a reduced size of the meta bit map at the expense of sometimes saving invalid data to the save volume. For example, FIG. 46 shows the interpretation of a meta bit map 691' having a granularity of two blocks per bit. Each bit is set if any one of the two corresponding blocks is valid, or conversely, each bit is clear only if neither of the two corresponding blocks is valid. In this case, the block address can be converted to a bit address by an integer division by two, for example, by an arithmetic right shift of the block address by one bit position.

Read-only Remounting of the Production File System

[000227] If a replication session experiences connectivity or bandwidth difficulties during transfer of the delta sets from the primary site to the secondary site, it is possible for the save volume to run out of space. There are three options for handling this possibility. The first option is to terminate the replication session, so that a new file system copy would need to be performed to reestablish a replication session. The second option is to freeze the production file system until save volume space becomes available. This freeze would interrupt all user access to the production file system. The third option

1 is to remount the production file system as read-only until the save volume has emptied
2 enough to handle incoming changes to the production file system.

3 4 Virtual Servers and Virtual Data Movers

5 [000228] There have been described above various facilities for replicating file
6 systems using the Internet Protocol. These facilities permit file systems to be replicated
7 from an active site to a disaster recovery site in an IP data network. In the event of a
8 disaster at the active site, the remote copy of the file system can be accessed at the
9 disaster recovery site.

10 [000229] It is desired to provide disaster recovery of not only file systems but
11 also the file servers that provide user access to the file systems. Aside from various
12 attributes in a file system, a file server stores a good deal of additional information
13 needed for enabling users to access the file system in accordance with a network file
14 access protocol. For example, this additional information includes security credentials
15 for authenticating the identity of users having permission to access the files in the file
16 system, and mounting information for finding files on removable storage media or remote
17 drives. By replicating this additional information from an active site to a disaster
18 recovery site, it is possible to resume user access at the disaster recovery site in
19 accordance with a network file access protocol after user access has been interrupted by a
20 disaster at the active site.

21 [000230] For replication of file servers, it is desirable to separate the software of
22 the file server into a set of files that define the functionality and user environment of the
23 file server independent of the hardware of the file server, and a set of programs that are

1 dependent on the hardware of the file server. The files that define the functionality and
2 user environment of the file server independent of the hardware of the file server
3 comprise a virtual file server. This virtual file server can be replicated over an IP
4 network from an active site to a disaster recovery site by replication of all of the files of
5 the virtual file server, and there is no need for the disaster recovery site to use the same
6 server hardware as the active site.

7 [000231] In many cases, a single data mover computer of a network file server
8 will contain a plurality of virtual file servers, and some configuration information will be
9 the same for these virtual file servers. In this case, it is desirable for the virtual file
10 servers to share certain files containing configuration information that is the same for the
11 plurality of virtual file servers. The collection of shared files and the other files of the
12 virtual file servers comprise a virtual data mover. The virtual data mover can be
13 replicated from an active site to a disaster recovery site by replicating all of the files of
14 the virtual data mover from the active site to the disaster recovery site.

15 [000232] FIG. 47 shows virtual data movers and virtual servers in the network
16 file server introduced in FIG. 1. The data mover 226 contains a virtual data mover 711
17 including a virtual server "A" 714 and a virtual server "B" 715. The data mover 227
18 contains a virtual data mover 712 including a virtual server "C" 716 and virtual server
19 "D" 717. The data mover 228 contains a virtual data mover 713 including a virtual
20 server "E". The cached disk array 229 contains the respective file systems 719, 720, 721,
21 722, and 723 of the virtual servers.

22 [000233] Virtual servers have a number of advantages unrelated to disaster
23 recovery. It is possible to replace a number of small file servers with a larger network

1 file server as shown in FIG. 47 containing a respective virtual server for each of the small
2 file servers. Each small file server can be replaced by a rather simple migration of the
3 server's files into a respective virtual server. The virtual servers become consolidated
4 into the larger network file server, yet each virtual server remains sufficiently isolated to
5 retain its original identity with respect to the network administrator and network users.

6 [000234] In a network file server having multiple data mover computers, the
7 virtual servers can be moved between the data mover computers to perform load
8 balancing. A virtual server can be moved by simply unmounting a file system of all the
9 virtual server's files from one data mover and mounting it on another data mover. It is
10 also very easy to migrate a virtual server to a remote site by taking a snapshot copy of
11 this file system and transmitting the snapshot copy to a remote site.

12 [000235] FIG. 48 shows various components of a data mover 725 prior to
13 creation of a virtual data mover. The data mover 725 services a group of Network
14 Information Service (NIS) and Domain Name Service (DNS) clients specified in a client
15 database 726. The data mover 725 has an NFS server 727 and a number of CIFS servers
16 728. The NFS server 727 and the CIFS servers 728 share a plurality of network
17 interfaces. Each of the servers has a global view of the resources in the data mover 725.

18 [000236] The data mover 725 has a root file system (FS) 729 shown in FIG. 48
19 to use a Unix file naming convention. All of the configuration and control data is stored
20 in the root file system 729 of the data mover. In addition, all the event logs are global to
21 all servers and are also stored in the root file system 72 of the data mover. The
22 /etc/(config) files include password, user group, virus checker configuration, Windows
23 Group Policy Object (GPO), and Kerberos files. Kerberos is a standard authentication

1 protocol characterized by mutual authentication of the client and server. The /.etc./(cifs
2 databases) include local group database, home directory files, and shares and registry
3 directories. The user file path names begin with “/fs” designating user file systems.

4 [000237] FIG. 49 shows that two virtual data movers 921, 922 have been
5 installed in the data mover 725 of FIG. 48. Each VDM 921, 922 contains a group of
6 CIFS servers working on a particular file system. Each VDM 921, 922 establishes a
7 logical link between a subset of the network interfaces and a file system of the data
8 mover that stores the CIFS databases for the VDM. For example, the virtual data mover
9 “VDM01” 921 establishes a logical link between the network interfaces 925 and the file
10 system 923 named “root_fs_vdm_1”, and the virtual data mover “VDM02” 922
11 establishes a logical link between the network interfaces 926 and the file system
12 “root_fs_vdm_2”. As shown, the root file system for the VDM is mounted onto a “/.etc”
13 file for the VDM in the root file system 729 of the physical data mover 725 containing
14 the VDM. The CIFS server is then able to retrieve the CIFS information from the client
15 connection that is using the data mover’s interface. This information includes an audit
16 log, local group and shares database stored in a set of files under the .etc private directory
17 of the file system associated with the VDM.

18 [000238] The virtual data movers 921, 922 provide isolation between their
19 servers as is needed in some environments such as Internet Service Providers (ISPs). The
20 virtual data movers also provide isolation between their respective file systems 923, 924,
21 which is useful for isolating data from disjoint user groups or departments hosted on the
22 same data mover.

1 **[000239]** A file system in a VDM can be exported via NFS in a fashion similar
2 to CIFS, so that the VDM NFS export will be relative to the VDM root. For example, to
3 export a file system “/fs11” in the VDM01, one may use a command such as
4 “server_export vdm_1 /fs11”.

5 **[000240]** FIG. 50 shows various components of a virtual data mover 730. These
6 components become stored in a root file system 747 for the virtual data mover 730, and
7 server file systems 746 for the virtual servers 731, 732 become mounted to the root file
8 system 747.

9 **[000241]** There are four types of components of the virtual data mover 730.
10 These types of components are: (1) items that can be contained entirely within the virtual
11 data mover; (2) those items which need to be enumerated and compared to a global set;
12 (3) those items which need to be “reserved or allocated” because they are physical
13 devices; and (4) those items which are specific to a site and its related architecture.

14 **[000242]** Items of the first type (fully contained within a VDM) include NFS
15 virtual servers 731, CIFS virtual servers 732, NFS exports 733, shares 734, Kerberos files
16 735, event logging 736, user mapping 737, password and group files 738, home directory
17 739, and NIS/DNS resolver items 740. Items of the second type (enumerated and
18 compared, and merged where able) include data mover parameters such as parameter for
19 setting the lookup technique for CIFS/NFS mapping (AD Mapping, Resolver,
20 useUnixGID, etc.). Items of the third type (tracked and reserved) include mounts and
21 interfaces. Items of the fourth type (global items affected by the “site”) include virus
22 checking parameters and time.

23 **[000243]** Mounts should be tracked inside of the virtual data mover, and when a

1 virtual data mover is installed, it should be queriable for the mounts that it expects. The
2 virtual data mover should be able to automatically mount the file systems that it needs if
3 they are available for mount on the host data mover. If there are conflicts, or unavailable
4 resources, then the load should either completely error out (in the case of a conflict) or
5 reflect that the mount was only partially successful (in the case of unavailability).

6 [000244] The interfaces expected by a virtual data mover should also be
7 queriable from the virtual data mover. For example, a request for information on a
8 specified virtual data mover should return something like:

9
10 Virtual DataMover Name: vCIFS01

11 Currently mounted on: <unmounted>

12 Uses Interfaces: cge0v1, cge0v2, cge0v3

13 CIFS servers: vCIFSa, vCIFSb, vCIFSd

14 Filesystems: /fs01 (mta_fs_01), /fs02 (mta_fs_02), /fs03 (mta_fs_03)

15
16 [000245] A virtual data mover should be considered a targetable object for
17 server level commands (e.g., log, mount, uptime, etc.). A virtual data mover should be
18 manageable in a similar fashion as a physical data mover. In addition, a virtual data
19 mover is a top-level object, like a file system, that can be mounted to a physical data
20 mover, migrated, or replicated. The mounting process will be called loading the virtual
21 data mover onto a physical data mover. Once loaded, a virtual data mover appears much
22 like another server. More than one virtual data mover could be loaded onto a single
23 physical data mover. For load balancing, a virtual data mover could be unloaded from

1 one physical data mover and loaded onto another physical data mover within the same
2 network file server. For migration, a virtual data mover could be unloaded from one
3 network file server and loaded on another network file server. For disaster protection, a
4 virtual data mover could be replicated from one network file server to another network
5 file server.

6 [000246] FIG. 51 shows a procedure for configuring a virtual data mover. Some
7 of these steps apply only to virtual data movers with a CIFS server, and other steps apply
8 to virtual data movers with a CIFS server or a NFS server. In step 751, the system
9 administrator configures the data mover interfaces (CIFS/NFS). In step 752, the system
10 administrator creates file systems and mounts them (CIFS/NFS). In step 753, the system
11 administrator sets data mover parameters (CIFS/NFS), unless default data mover
12 parameters are to be used. In step 754, the system administrator edits Kerberos files to
13 join servers for secure communication (NFS/CIFS). In step 755, the system administrator
14 creates exports and shares (CIFS/NFS). In step 756, the system administrator edits user
15 and user group password files (CIFS/NFS), if static mappings are used. In step 757, the
16 system administrator configures and runs a user mapper (CIFS) in order to set user access
17 privileges, if automatic mapping is to be used. The user mapper is a service that assigns
18 UNIX style user identifiers (UIDs) and user group identifiers (GIDs) to CIFS clients. In
19 step 758, the system administrator configures quotas (CIFS/NFS), if quotas are required.
20 The quotas specify limits on storage capacity available to specified users or user groups.
21 In step 759, the system administrator configures home directories (CIFS) if applicable.
22 In step 760, the system administrator edits the Windows registry (CIFS) if needed.

23 [000247] Once a virtual data mover has been configured, no additional

1 configuration should be required to fully move the entire configuration from one physical
2 data mover to another. All of the required configuration information should be contained
3 within the configuration files depending from the root directory of the virtual data mover.

4 **[000248]** Virus checking is configured at the site level. If virus checking is to
5 be run at the secondary site, it should be configured there independently. When virus
6 checking is run at the secondary site, it is recommended that a scan should be initiated on
7 a first read after failover. This will ensure scanning of any virus infected files that were
8 on the virus checking queue at the time that the production file was paused to take a delta
9 set for transmission to the secondary site. Alternatively, the virus checking process at the
10 primary could be coordinated with delta set creation so that only virus-free delta sets are
11 created. For example, the normal pause during delta set creation could be extended to
12 flush the virus checking queue before the delta set is taken. In this fashion, the integrity
13 of the file system with respect to virus checking could be preserved during a failover and
14 failback, as well as more typical operations such as migration and backup.

15 **[000249]** A graphical user interface (GUI) has been used for management of
16 physical data movers within a network file server. This GUI is extended to provide
17 management of virtual data movers within the network file server or hosted by a selected
18 physical data mover in the network file server. FIG. 52, for example, shows various
19 virtual data mover management functions that are selectable from the GUI. These
20 functions include create a new VDM (771), list VDMs (772), select a VDM (773), load a
21 selected VDM (774), unload a selected VDM (775), list virtual servers for a selected
22 VDM (776), select a virtual server (777), query information for a selected VDM (778),
23 query dependencies for a selected VDM (779), estimate interfaces for a selected VDM

1 (780), estimate mounts for a selected VDM (781), query information for a selected virtual
2 server (782), mount file systems for a selected VDM (783), replicate a selected VDM
3 (784), failover a selected VDM (785), failback a selected VDM (786), resync a selected
4 VDM (787), and replicate a selected virtual server (788).

5 **[000250]** Creation of a new VDM (771) involves allocation of a new root file
6 system of the VDM unless the user specifies an existing root file system. The default
7 size for a new root file system, for example, is 128 megabytes, which is the same as the
8 default size for a new root file system for a physical data mover. On success, the control
9 station of the network file server saves the binding relationship between the VDM and the
10 root file system into the control station database for further management.

11 **[000251]** Query dependencies for a selected VDM (779) is often executed
12 before a configuration change on the primary. For instance, dependencies on a selected
13 VDM can be checked to ensure that the interfaces or mounts on the primary also exist on
14 the secondary. It is also possible to check for dependencies for a proposed configuration
15 change such as addition of a CIFS server on the primary. The system administrator will
16 be advised of any interfaces or mounts for the proposed CIFS server that are missing
17 from the primary or secondary.

18 **[000252]** A load of a selected VDM (774) onto a physical data mover involves
19 the physical data mover checking whether the VDM configuration information (from the
20 procedure of FIG. 51) has been stored on the VDM root file system, and if not, the
21 physical data mover will require creation of the VDM configuration information. Once
22 the configuration information has been stored on the VDM root file system, the physical
23 data mover retrieves the configuration information from the root file system and stores it

1 in data mover memory. However, the load operation will fail if the VDM is already
2 loaded onto any physical data mover in the network file server.

3 **[000253]** When a VDM is loaded, the physical data mover plays a CIFS setup
4 stored in a `/etc/vdm.cfg` file to retrieve the CIFS configuration of the VDM. During this
5 operation, the interface assignment is applied to each CIFS server, and warning messages
6 may be reported in the server's log if a bad interface (already in used or unknown) is
7 specified. In that case, the administrator will have to change the CIFS configuration to
8 setup the right interface configuration.

9 **[000254]** A `server_mountpoint` command is provided to permit a system
10 administrator to create a mountpoint on a specified VDM, and to check the mountpoints
11 that are currently created on the VDM. These mountpoints are relative to the VDM. For
12 example, a mountpoint `/fs1` is actually `/root_fs_vdm1/fs` from the data mover point of
13 view. Commands are also provided for mounting a file system and unmounting a file
14 system. For a raw file system, the mountpoint is required to be created using the
15 `server_mountpoint` command before a server mount operation. For a Unix-based file
16 system, if a mountpoint is not specified, it will be automatically created when mounting a
17 file system on a VDM and cleaned up when unmounted. The default file system name
18 will be the same as the file system name; e.g., a file system named "fs" will be mounted
19 as `/root_fs_vdm1/fs`. A system administrator can only mount or unmount a file system on
20 a VDM when the VDM has been loaded. The control station will direct mount and
21 unmount commands to the physical data mover to which the VDM has been loaded. For
22 example, the system administrator can load a VDM in a passive mode, mount file
23 systems to the VDM, and then activate the VDM.

1 **[000255]** A `server_cifs` command is provided to allow configuring a WINS
2 configuration for a VDM. If no particular settings are applied to a VDM, the settings are
3 defaulted to the values of the physical data mover containing the VDM.

4 **[000256]** Any user mapper configuration, which is for an external user mapper,
5 is global to the physical data mover just as NIS and the password file. In general, for a
6 CIFS client to access a UNIX-based file system, the CIFS client must be assigned a
7 unique UNIX style user identifier (UID) and a corresponding user group identifier (GID).
8 This enables the CIFS client to access the UNIX file permissions and quotas. The
9 mappings of CIFS clients to UNIX style UIDs and GIDs can be established in a number
10 of different ways, such as storage of the mappings in a static file of mappings, storage of
11 the mappings in the Microsoft Active Directory (AD), storage of the mappings in NIS, or
12 by running the user mapper service.

13 **[000257]** The user mapper service runs on the data mover, and it automatically
14 and persistently allocates UIDs and GIDs from a range set aside for CIFS users.
15 Therefore, the user mapper service simplifies CIFS only environments by taking all
16 administration out of managing these mappings. When a CIFS user establishes an initial
17 connection with a data mover or a VDM, the data mover or VDM requests the mapping
18 from the user mapper, and the mapping from the user mapper is persistently cached in the
19 root directory of the data mover or VDM. After the user mapping is cached, the data
20 mover or VDM need not request the user mapping for the user from the user mapper,
21 because the data mover or VDM may obtain the user mapping from the root directory of
22 the data mover or VDM.

1 **[000258]** Stopping and starting CIFS threads (which process all VDM and non-
2 VDM CIFS requests) is global to the physical data mover. A `server_setup server_x -P`
3 `cifs -o stop/start` command affects all the VDMs of the physical data mover named
4 “server_x”.

5 **[000259]** Having a CIFS configuration file per VDM allows use of a different
6 CIFS setup for WINS and home directory strategy on the same physical data mover.
7 The share and local group databases (e.g., Kerberos and other configuration databases)
8 are split per VDM basis. The share paths are now relative to the VDM’s file system,
9 which allows abstracting of the mount point path of the file system.

10 **[000260]** When a VDM is unloaded, the CIFS servers of the VDM are stopped
11 and removed from the Data Mover, including release of DNS and WINS host
12 registrations. If a computer name is defined in the VDM configuration, the unjoin from
13 the Windows domain should be issued before unloading the VDM.

14 **[000261]** A “`server_cifs -delete`” command is provided to remove one or more
15 interfaces from a CIFS server and to put them in a list of unused VDM interfaces.

16 **[000262]** A VDM move command is provided to combine an unload of a
17 selected VDM from one physical data mover and a load of the selected VDM on another
18 physical data mover. The move command also internally takes care of the unmount and
19 mount of the non-root virtual server file systems contained within the VDM, so that the
20 virtual servers contained within the selected VDM are moved as well.

21 **[000263]** A VDM can be moved from one physical data mover to another
22 without changing the share configuration and no local group migration is required. For
23 example:

1

2 `nas_server -vdm vdm_1 -move server_2`

3

4 The `nas_server` command can specify interfaces to assign to the VDM. If one of these
5 interfaces is already used by another VDM, the command fails. So, all the specified
6 interfaces should be in the unused VDM interface list to complete the command
7 successfully.

8 **[000264]** A move command also can be used to move an existing CIFS server of
9 the “global” CIFS servers located on the root file system of the physical data mover to the
10 VDM loaded on top of it or move it back. For example, in response to the command:

11

12 `server_cifs <source_mover> -move netbios/comptype=<server_name> <target_mover>`

13

14 the interfaces connected to the CIFS server are moved from the source server to the target
15 server. The share and local group databases for the CIFS server are moved to the private
16 directory of the target server. The `vdm.cfg/boot.cfg` files of the source and destination
17 servers are updated so that the physical data mover is able to reload the CIFS
18 configuration at boot time.

19 **[000265]** Kerberos information can be split into each VDM so that each VDM
20 keeps its own keytab and configuration files. In this case CIFS servers are initially joined
21 to the VDM, and credentials need not be moved into and out of VDMs. This means that
22 CIFS servers are initially created in their VDMs, and joined to the domain. In this way,
23 their kerberos information is stored in the VDM root file system. In this manner, a VDM

1 can be unloaded from one physical data mover and loaded onto another, without the need
2 to rejoin the domain.

3 [000266] A VDM delete operation is provided to allow a system administrator
4 to irrevocably delete a VDM including the entire CIFS configuration (local group and
5 share databases) and the VDM root file system. All non-root file systems mounted on
6 the VDM are required to be unmounted before the deletion.

7 [000267] A VDM access control level command is provided to set an access
8 control level value that defines the owner of a VDM, or the level of access allowed for
9 users and groups defined in an access control level table for the VDM.

10 [000268] A VDM replicate command takes a VDM as an argument, and by
11 default replicates, as a group, the VDM root file system as well as the data file systems
12 associated with that VDM. The system administrator may also choose partial replication
13 of the VDM file systems. The system administrator may request information about the
14 status of replication. For example, in response to a “replicate info” command, the file
15 server returns something like:

16
17 <state> <rate> <rep_state> <#deltas>
18 Virtual DataMover Name: vCIFS01
19 Currently mounted on: server_2
20 Filesystems: vdm_vCIFS01 replicated 700k/s alive 1
21 mta_fs_01 replicated 2000k/s alive 3
22 mta_fs_02 non-replicated n/a n/a n/a
23 mta_fs_03 replicated n/a out-of-sync n/a

[000269] If the replication information indicates that a file system has fallen out of sync (or has never been replicated), then the system administrator can individually reinitialize the replication session.

[000270] A VDM failover command results in the failover of the file systems (VDM root and supporting file systems) as well as unloading the VDM on the primary file server and loading of the VDM on the secondary file server. In a similar fashion, a VDM failback command results in the failback of the file systems (VDM root and supporting file systems) as well as unloading the VDM on the secondary file server and loading of the VDM on the primary file server. Failover and failback, however, should be performed on an entire VDM and all of its available supporting file systems, and not just one supporting file system. Replication should support quality of service (e.g., network bandwidth throttling) at the file system level, as well as the ability to consistently snapshot a VDM and its file systems. There should be an option to consistently coordinate the delta set creation and playback for a set of file systems. Group consistent snapshots should also be supported on the read-only secondary file system side.

[000271] Failover and failback semantics should support synchronously or asynchronously issuing a failover or a failback command. In the case of a synchronous failover, a failover command is issued, production file systems are mounted read-only, final production side delta sets are cut, final production side delta sets are replayed on the secondary, and then the secondary file systems are mounted read-write. This provides

1 the system administrator with a way to do a failover of a VDM with no data loss. For the
2 case of asynchronous failover, a failover command is issued for the VDM, all available
3 delta sets are replayed on the secondary side, the production file systems are mounted
4 read-only (or ignored, depending on whether the primary is available), and the secondary
5 file systems are taken read-write. The failover commands (especially for the
6 synchronous case), should be executed in parallel by the control station, and not
7 sequentially.

8 [000272] It should be possible for a system administrator to manage the virtual
9 data movers as shown in FIG. 52 from a web-based GUI provided by the control station
10 (230 in FIG. 47) of the network file server. At the control station level, the VDMs are
11 manageable server objects that can be loaded and unloaded, and operated on from the
12 replication level. The web-based GUI should permit the system administrator to create a
13 new VDM, load a VDM, unload a VDM, query information about a VDM, query
14 dependencies for a VDM, list VDMs, enumerate the expected interfaces, enumerate the
15 expected mounts, and automatically mount file systems as far as possible. Once a VDM
16 is loaded, the system administrator should be permitted to manage shares, CIFS servers,
17 file systems, and the other aspects of the VDM configuration. The system administrator
18 should be permitted to list the available VDMs, choose one, and replicate the VDM and
19 some of the data file systems associated with the VDM. The system administrator should
20 be permitted to choose a VDM and failover from that VDM, resync that VDM, and
21 failback to that VDM.

22 [000273] The web-based GUI should permit the system administrator to monitor
23 replication initialization and steady state from a top-level summary perspective. For

1 example, the GUI should show all file system replications for a specified virtual server in
2 a specified VDM. For replication of a specified virtual server, the GUI should show the
3 number of delta sets outstanding and the state of replication and playback. The GUI
4 should permit the system administrator to set events based on a specified percentage of
5 the save volume being full or a specified number of delta sets outstanding for alerting the
6 operator of impending normal or abnormal completion of the replication process.

7 [000274] In general, for replication of a VDM from a primary site to a
8 secondary site, it is not necessary for there to be any checking or reservation of resources
9 on the secondary site that would be needed for loading of the VDM at the secondary site.
10 It is recommended, however, that items such as mounts and interfaces (as well as non-
11 VDM aware configuration items such as DNS, NIS, time, and parameters) should be
12 validated to ensure that mounts are identical on the primary file server and the secondary
13 file server, interfaces are identical and functioning on the secondary file server, and the
14 DNS, time, and parameters are properly configured on the secondary file server. Once a
15 disaster recovery environment is established, if additional interfaces or file systems are
16 added to the VDM at the primary file server (and the changes to the VDM are replicated
17 to the secondary file server), then the secondary file server will automatically be checked
18 for the additional interfaces or file systems.

19 [000275] During loading or operation of a VDM at the secondary site after
20 failover, missing items or improperly configured parameters may cause problems.
21 Therefore, the system administrator may issue a command for comparing the
22 requirements for interfaces and mounts of a VDM with the resources available on a
23 secondary file server. In response to the command, the secondary file server provides the

1 system administrator with a list of any discrepancies that should normally prevent a
2 VDM from loading on the secondary file system. Such discrepancies would include
3 missing file systems or file systems that cannot be mounted, insufficient interfaces, or
4 inconsistent configuration parameters. For example, a file system can be mounted in
5 only one VDM at any given time. The system administrator may issue this command on
6 demand, to verify the disaster recovery environment after a change. The system
7 administrator may also run the command prior to failover.

8 9 Prevalidation and Resource Reservation for a Configuration Change or for 10 Failover

11 [000276] In a typical case, it is desirable to perform prevalidation and resource
12 reservation when a configuration change is made to the primary that may need additional
13 resources such as interfaces or mounts on the secondary side. It is also desirable to
14 perform prevalidation and resource reservation just prior to failover of a VDM, in case
15 the system has been operating without sufficient resources on the secondary side to
16 ensure full user access at the disaster recovery site after failover from the active site to the
17 disaster recovery site. FIG. 53 shows such a procedure, which performs prevalidation,
18 resource reservation, and then a configuration change or failover. The procedure, for
19 example, is performed by the control station of the primary file server in response to a
20 configuration change request or a failover request from a system administrator. The
21 procedure could also be performed by the control station of the secondary file server in
22 response to a disaster at the primary file server, for a case where the VDM has already
23 been replicated to the secondary file server.

1 **[000277]** In a first step 791, the control station determines the requirements for
2 interfaces and mounts for the VDM on the primary site. For a desired configuration
3 change, the requirements include the interfaces and mounts needed for the VDM to
4 support the desired configuration change. Then in step 792, the control station
5 determines the resources available on the secondary site for interfaces and mounts for the
6 VDM. In step 793, if there are sufficient resources on the secondary site for the proper
7 loading and operation of the VDM on the secondary site to support any desired
8 configuration change and to support failover, then execution branches to step 794. In
9 step 794, the control station reserves all of the resources on the secondary site for the
10 interfaces and mounts needed to support any desired configuration change and for
11 failover of the VDM to the secondary site. Then, in step 795, the control station performs
12 the desired configuration change or the failover of the VDM from the primary site to the
13 secondary site, and the procedure is finished.

14 **[000278]** In step 793, if there are insufficient resources on the secondary site for
15 all of the required interfaces and mounts of the VDM, then execution continues to step
16 796. In step 796, the control station provides the system administrator with a list of the
17 missing resources or discrepancies that may cause problems with the loading or operation
18 of the VDM on the secondary file server if any desired configuration change is made and
19 failover would occur. In step 797, the control station receives the system administrator's
20 choice of whether to terminate or to force a configuration change or to force failover. If
21 the system administrator chooses termination, then the procedure is finished. Otherwise,
22 if the system administrator chooses to force a configuration change or to force failover,
23 then in step 799 the control station reserves available resources on the secondary site

1 needed for interfaces and mounts if any desired configuration change is made and
2 failover of the VDM would occur. Execution continues from step 788 to step 795, to
3 perform the desired configuration change or failover of the VDM from the primary file
4 server to the secondary file server. After step 795, the procedure is finished.

6 Continuous Availability of User Mappings During Failover

7 [000279] In a typical disaster recovery environment, users log onto a virtual
8 server in a virtual data mover on a primary file server at an active site. During user
9 access, changes to the root file system of the virtual data mover and the mounted file
10 systems of the virtual servers are replicated to a secondary file server at a disaster
11 recovery site. When a disaster occurs at the active site, the virtual data movers are
12 activated at the disaster recovery site in a failover operation. The users can then log onto
13 the virtual servers at the disaster recovery site, and continue data access operations.

14 [000280] For some critical applications, the network traffic and processing
15 resources for enabling a user to log onto a server are problematic in the disaster recovery
16 environment. This is especially true for a CIFS server. For example, a typical user
17 logging onto a Windows 2000 server may use eight different protocols (SMB, ICMP,
18 UDP, NBT, TCP, MSRPC, LDAP, and Kerberos) transferring about 31 kilobytes of data
19 over the network in 167 frames. See, for example, Greg Molnar et al., "Windows 2000
20 Startup and Logon Traffic Analysis," Microsoft Enterprises Services, White Paper, Aug.
21 2000, pp. 1-34, Microsoft Corp., www.microsoft.com. Continuance of user access by the
22 user logging into the secondary server in the event of a disaster is impeded by the need
23 for network resources at a time when network resources are least available.

1 **[000281]** One way of providing continuance of user access in the event of a
2 disaster is to provide mirroring of user mapping information from the logon process from
3 the active site to the disaster recovery site. This provides continuance of user access
4 unless the disaster occurs at the active site before all of the user mapping information
5 from the logon process is transmitted from the active site.

6 **[000282]** FIG. 54 shows an example of a disaster recovery environment in
7 which a secondary data mover 801 at a disaster recovery site handles user mapping for
8 the primary data mover 730 at an active site 730. A primary copy of user mappings 807
9 is stored in the root of the secondary root file system 806. The user mappings permit
10 current users of the primary copy of the VDM file system 747 (and selected primary
11 server file systems 746) to continue to access the secondary copy of the VDM file system
12 806 (and selected secondary server file systems 805).

13 **[000283]** The virtual data mover 713 contained in the primary data mover 730
14 has a user mapper relay 809 for intercepting user mapping requests. The user mapper
15 relay 809 forwards to a user mapper primary 808 in the secondary virtual data mover 802
16 any request for user mapping information that is not found in a read-only user mapping
17 cache 810 in the primary root file system 747. For example, the user mapper relay 809
18 forwards to the user mapper primary 808 a request for user mapping information of a new
19 user that is logging into the primary virtual data mover 713. The user mapper primary
20 808 accesses the primary user mappings 807 and returns the user mappings for the new
21 user to the user mapper relay 809. The user mapper relay 809 stores the user mapping
22 information for the new user in the read-only user mapping cache 810 in the root of the
23 primary root file system 747.

1 **[000284]** In the case of a disaster at the disaster recovery site 801, the user
2 mappings for the active users are found in the user mapping cache 810. The user mapper
3 relay 809 is then promoted to primary status and takes over the functions of the user
4 mapper primary 806 and stores any new user mappings in the user mapping cache 810.
5 Once operation at the disaster recovery site 801 is restored, the new user mappings are
6 migrated from the user mapping cache 810 to the primary user mappings 807 at the
7 disaster recovery site 801, and then the user mapper primary 808 returns to primary status
8 and the user mapper relay 809 returns to secondary status.

9 **[000285]** FIG. 55 shows a procedure for using the user mapping facilities of
10 FIG. 54. In a first step 811, the user sends a logon request to the virtual server at the
11 active site. In step 812, user mapper relay in the primary data mover checks the user
12 mapping cache at the active site for mapping information for the user. If the user
13 mapping cache does not have current mapping information for the user, then the user is a
14 new user and execution continues from step 812 to step 813. In step 813, the user
15 mapper relay in the primary data mover sends a request for the user's mapping
16 information to the user mapper primary in the secondary data mover at the disaster
17 recovery site. In step 814, the user mapper primary in the secondary data mover accesses
18 the user's user mapping information in the primary user mappings and returns the user's
19 user mapping information to the user mapper relay. Finally, in step 814, the user mapper
20 relay caches the user's user mapping information in the user mapping cache and begins
21 the session with the user.

1 **[000286]** In step 812, if the user is not a new user, then execution branches from
2 step 812 to step 816. In step 816, the user mapper relay reads the user's user mapping
3 information from the user mapping cache and begins the session with the user.
4
5

6 Elimination of CIFS Disconnection During Failover

7 **[000287]** A further enhancement to the user mapping facilities in the disaster
8 recovery environment of FIG. 54 would be the elimination of CIFS disconnections in the
9 event of failover. CIFS disconnections may result when requested access to a file system
10 is denied for longer than a certain timeout interval. For example, CIFS session
11 disconnects may result from freeze/thaw operations that normally occur with the
12 playback of delta sets during the replication process if the freeze/thaw operation is not
13 completed within the timeout interval. These CIFS session disconnects could be
14 eliminated by the CIFS interface in the data mover setting a timer upon receipt of a CIFS
15 request in order to always return a response to the CIFS client within the timeout interval.

16 **[000288]** In the disaster recovery environment of FIG. 54, another possible
17 cause of CIFS disconnects would be a disaster at the active site causing a disconnection
18 of the CIFS session between a CIFS client and the primary data mover 730. These CIFS
19 session disconnects could be eliminated in the disaster recovery environment of FIG. 54
20 by reprogramming the CIFS client to direct further CIFS requests to the disaster recovery
21 site if the active site 730 fails to respond with the timeout interval. The secondary data
22 mover 801 would recognize that the CIFS client is an active client, inspect the health of
23 the active site, and initiate a failover from the active site to the disaster recovery site if the

1 active site would be inoperative. The primary data mover could confirm to the CIFS
2 client that further CIFS requests should be sent to the disaster recovery site, or else the IP
3 network interconnecting the CIFS client, the active site, and the disaster recovery site
4 could be re-programmed so that the disaster recovery site 801 would assume the network
5 identify of the active site and would receive any further CIFS requests addressed to the
6 active site.

7 [000289] FIG. 56 shows a specific example of a disaster recovery environment
8 for maintaining CIFS connections during a disaster. An IP network 820 interconnects a
9 CIFS client 821 to a primary CIFS server 822 in the primary data mover 730 and to a
10 secondary CIFS server 829 in the secondary data mover 801. For maintaining CIFS
11 connections, a primary copy of user session information 811 is collected in the primary
12 data mover 730 and synchronously mirrored to a secondary copy 812 in the secondary
13 data mover 801. This synchronous mirroring of the user session information could be
14 performed by transmitting the user session information from the primary data mover 730
15 to the secondary data mover 801 over the IP network 820 or by transmitting the user
16 session information over a dedicated data transmission link from the primary data mover
17 730 to the secondary data mover 801.

18 [000290] The user session information 811, 812 identifies each user having an
19 active session with the primary data mover, and includes the credentials assigned to each
20 user during each active session. In this fashion, a user session may be interrupted at the
21 active site and resumed at the disaster recovery site using the credentials in the secondary
22 copy 812 of the user session information without a need for the user to login to the
23 secondary data mover at the disaster recovery site. The user session information 811,

812 may also include a sufficient amount of session state information that is synchronously mirrored from the primary data mover to the secondary data mover to avoid disruption of an active session when the active session is interrupted at the active site and resumed at the disaster recovery site. For example, an active session is subdivided into a series of transactions, and session state information that is committed to the primary copy 811 of the user session information at the end of each transaction is also committed to the secondary copy 812 of user session information before the session proceeds to the next transaction.

[000291] A CIFS connection maintenance agent 823 is interposed between the IP network 820 and the primary CIFS server 822 in order to intercept CIFS requests from a CIFS agent 825 in the CIFS client 821 and to pass the CIFS requests to the primary CIFS server. The CIFS connection maintenance agent 823 has a timer 824 for determining when the primary CIFS server 822 fails to respond to a CIFS request within a first time interval that is less than the timeout interval for a CIFS connection. When the primary CIFS server 822 fails to respond to a CIFS request within the first time interval, the CIFS connection maintenance agent 823 returns a “keep alive” response to the CIFS agent 825, and also logs the problem and reports the problem to the system administrator.

[000292] The CIFS client 821 has a CIFS redirection agent 826 interposed between the CIFS agent 825 and the IP network 820 for intercepting requests and responses between the CIFS agent 825 and the IP network. The CIFS redirection agent 826 has a timer 827 for determining when the primary data mover 730 fails to respond to a CIFS request within a second time interval that is greater than the first time interval and less than the CIFS connection timeout interval. In this case, the CIFS redirection agent

1 will return a “keep alive” response to the CIFS, and the CIFS agent will retransmit the
2 CIFS request. After a CIFS request has been retransmitted a certain number of times to
3 the primary data mover without a response, the CIFS redirection agent 826 accesses a
4 database 828 of alternative network addresses. Upon finding that the database 828
5 contains an alternative address of a secondary CIFS server 829 registered for the primary
6 CIFS server 822, the CIFS redirection agent 826 redirects the CIFS request to the
7 secondary CIFS server 829 in the secondary data mover 801.

8 [000293] The secondary data mover 801 has a CIFS connection maintenance
9 agent 830 interposed between the IP network 820 and the secondary CIFS server 829 for
10 intercepting redirected CIFS requests from network clients. The CIFS connection
11 maintenance agent 830 reports any redirected CIFS request to the control station 830 for
12 the secondary data mover 801. The control station determines whether or not to initiate
13 failover from the primary CIFS server to the secondary CIFS server. Once failover has
14 been completed, the CIFS connection maintenance agent 830 passes the CIFS request to
15 the secondary CIFS server 829 in order to continue the CIFS session. Then the CIFS
16 connection maintenance agent functions in a fashion similar to the CIFS connection
17 maintenance agent 823 in the primary data mover. In other words, the CIFS connection
18 maintenance agent 830 operates a timer 832 to determine when the secondary CIFS
19 server 829 fails to respond to a CIFS request within the first time interval, and if so, to
20 return a “keep alive” response to the CIFS client 821 in order to maintain the CIFS
21 connection.

22 [000294] FIG. 57 summarizes the operation of the CIFS connection and
23 redirection agents in the disaster recovery environment of FIG. 56. In a first step 841, the

1 CIFS connection maintenance agent in the primary data mover normally provides a
2 response to a CIFS request with the first time interval, which is less than the CIFS
3 disconnection time-out interval. Next, in step 842, the CIFS redirection agent in the
4 CIFS client detects a failure of the CIFS primary server in the primary data mover to
5 respond within the first time interval and inhibits disconnection and redirects the CIFS
6 request to an alternative network address of a secondary CIFS server in the secondary
7 data mover. Finally, in step 843, the CIFS connection maintenance agent in the
8 secondary data mover receives the redirected CIFS request, notifies its control station to
9 initiate failover of the primary CIFS server to the secondary CIFS server, and continues
10 the CIFS session with the secondary CIFS server.

11 [000295] The discussion above concerning the replication of virtual data movers
12 from an active site to a data recovery site has focused on techniques for reducing or
13 eliminating disruption due to a disaster at the active site or loss of communication
14 between the user and the active site. The replication of virtual data movers, however, can
15 be particularly useful for various situations where the disruption cannot be eliminated or
16 when a disruption does not occur. One such situation is where it is desired to maintain a
17 series of snapshot copies of the virtual data mover for recovery from software and data
18 corruption, and for data mining tasks. In this situation, the snapshot copies of the virtual
19 data mover can be created at the disaster recovery site without use of any processing
20 capability or storage capacity at the primary site.

21 [000296] FIG. 58, for example, shows that the secondary data mover 801 has
22 created a series of snapshot copies of the virtual data mover 802 at respective points in
23 time at $t=t_0$, $t=t_1$, and $t=t_2$. It is desired for each of these snapshot copies to have a

1 consistent state for the virtual data mover, and it is also desired for the user files to be in a
2 group consistent state in the user file systems (FS1, FS2) in each snapshot copy 851, 852,
3 853. For example, if the user files include financial accounts in which money is
4 transferred from an account in the first user file system FS1 to an account in the second
5 user file system FS2, the user file systems would be group consistent when the accounts
6 balance. The accounts would balance before the transaction and once the transaction has
7 completed. Therefore, it would be possible to permit user access to resume from any
8 such prior consistent state of the virtual data mover.

9 [000297] For example, if the primary data mover 730 were to crash and it were
10 found that the files associated with the primary data mover were corrupted, an attempt
11 would be made to restore the virtual data mover 712 at the primary site with the present
12 state of the virtual data mover 802 at the disaster recovery site. However, if the
13 corruption were caused by a virus or by malicious tampering, it is possible that the
14 present state of the virtual data mover 802 could also be corrupted. Under these
15 circumstances, it would be desirable to restore the virtual data mover 713 at the active
16 site with the state of the most recent one of the snapshot copies 853, 852, 853 that has not
17 been corrupted.

18 [000298] FIG. 59 is a flowchart of a procedure used at the disaster site for
19 creating the series of snapshot copies of a secondary virtual data mover at the disaster
20 recovery site. Basically, a background process creates each snapshot copy concurrent
21 with the updating of the secondary virtual data mover with delta sets from the active site.
22 In a first step 861, the disaster recovery site receives a delta set from the active site. The
23 delta set is a consistent modification to the environment of the secondary virtual data

1 mover. In step 862, the disaster recovery site updates the secondary virtual data mover
2 with the delta set. In step 863, if it is not time for a new snapshot of the secondary virtual
3 data mover, then execution loops back to step 861.

4 [000299] Otherwise, if it is time for a new snapshot of the secondary virtual data
5 mover, then in step 864, the disaster recovery site initiates a process of making a snapshot
6 copy of the secondary virtual data mover. For example, snapshots for backup purposes
7 can be made periodically, and before or after a major configuration change, such as
8 before the deletion of a virtual server, or after the addition of a virtual server. Execution
9 loops from step 864 back to step 861. Most of the processing for making the snapshot
10 copy occurs in background, after execution loops from step 864 back to step 861.

11 [000300] In order for user files to be in a group consistent state when a snapshot
12 is made, user access at the active site can be temporarily suspended while current
13 transactions complete, and any modifications from these current transactions are included
14 in delta sets transmitted from the active site to the disaster recovery site. Once the
15 modifications from these current transactions are included in delta sets transmitted to the
16 disaster recovery site, user access at the active site can be resumed. Once the
17 modification from these current transactions have been received in delta sets transmitted
18 to the disaster recovery site and integrated into the secondary copies of the user files at
19 the disaster recovery site, a snapshot copy of the secondary virtual data mover is made at
20 the disaster recovery site.

21
22 [000301] In view of the above, there has been described a disaster recovery
23 system in which files that define the user environment of a file server at an active site are

1 replicated to produce a virtual server at a disaster recovery site remote from the active
2 site. To discontinue user access at the active site and resume user access at the disaster
3 recovery site, the disaster recovery system determines whether there are sufficient
4 network interfaces and file system mounts at the disaster recovery site for the virtual
5 server for providing user access at the disaster recovery site. If so, the required resources
6 are reserved for the virtual file server, and user access is switched over from the active
7 site to the disaster recovery site. If not, an operator is given a list of missing resources or
8 discrepancies, and a choice of termination or forced failover. Interruptions during the
9 failover can be avoided by maintaining a copy of user mappings and a copy of session
10 information at the disaster recovery site, and keeping alive client-server connections and
11 re-directing client requests from the active site to the disaster recovery site during a
12 failure at the active site.

13